

[HTTPS://SYSTEMSAPPROACH.ORG](https://systemsapproach.org)

Private 5G: A Systems Approach

Larry Peterson, Oguz Sunay,
and Bruce Davie

Table of Contents

- [Foreword](#)
- [Preface](#)

- Acknowledgements
- Chapter 1: Introduction
 - 1.1 Standardization Landscape
 - 1.2 Access Networks
 - 1.3 Managed Cloud Service
 - 1.4 Beyond 5G
- Chapter 2: Architecture
 - 2.1 Overview
 - 2.2 Radio Transmission
 - 2.3 Radio Access Network
 - 2.4 Mobile Core
 - 2.5 Managed Cloud Service
- Chapter 3: Radio Transmission
 - 3.1 Coding and Modulation
 - 3.2 Scheduler
 - 3.3 Virtualized Scheduler (Slicing)
 - 3.4 New Use Cases
- Chapter 4: Radio Access Network
 - 4.1 Packet Processing Pipeline
 - 4.2 Split RAN
 - 4.3 Software-Defined RAN
 - 4.4 Near Real-Time RIC
 - 4.5 Control Loops
- Chapter 5: Mobile Core
 - 5.1 Identity Management
 - 5.2 Functional Components
 - 5.3 Control Plane
 - 5.4 User Plane
- Chapter 6: Managed Cloud Service
 - 6.1 Building Blocks

- [6.2 Example Deployment](#)
- [6.3 Cloud Management Platform](#)
- [6.4 Connectivity API](#)

- [About The Book](#)

- [Read the Book](#)
- [Build the Book](#)
- [Contribute to the Book](#)

- [About The Authors](#)

- [*Read the Latest!*](#)

- [*Print Copies*](#)

Foreword

When I was a graduate student studying parallel computing in the early 1990s, the World Wide Web exploded onto the scene, and I wanted to understand how the Internet made such rapid advances possible. I picked up a copy of the first edition of *Computer Networks: A Systems Approach*, and I started reading. I was delighted to learn not only the Internet's history and main concepts, but also to see examples of real code that illustrated how the protocols actually worked and how they could be changed. I was hooked. I loved the idea of working in the computer networking field, where the technologies we build can help bring people together and lower the barriers to innovation so that anyone who can write software can contribute.

As time went on, I saw that, while the Internet enabled rapid advances in end-host applications, the inside of the Internet was harder to change. Network devices were closed and proprietary, and protocol standards evolved slowly and painfully. In response, enabling innovation inside the network became a passion of mine, and of many other technologists eager to make the Internet better—more secure, performant, reliable, cost-effective, and easier to manage. Gradually, the computer networking field changed, with the advent of software-defined networking, giving network owners—such as the hyperscalers running large data centers—much more control over the software inside their networks. (See *Software-Defined Networks: A Systems Approach* for more of this story!) Nowadays, computer networking really is about software. It's a welcome change.

In recent years, the excitement has moved to the network edge with the emergence of 5G cellular access networks. No longer are wireless communication, computer networking, and cloud

computing siloed parts of some larger ecosystem. They are brought together, allowing wireless devices to connect to nearby computing resources. Far beyond providing mobile Internet access, these networks enable exciting real-time applications like augmented and virtual reality (AR/VR), Internet of Things (IoT), self-driving cars, drones, robotic control, and more. Plus, these networks are not only the purview of large carriers that must invest significant resources in wireless spectrum, cell towers, and more. Rather, individual enterprises, campuses, and communities are deploying private 5G within their own organizations to support their own applications, using lightly regulated spectrum like CBRS (Citizens Broadband Radio Service). Opportunities for innovation abound!

However, now there is so much more to learn, to understand the many parts of 5G access networks as a single coherent system. The learning curve can be steep. Even fairly technical people usually know one or at most two parts of the system, and do not know how the parts relate to the larger whole. For example, I knew a good amount about computer networking, a little about cloud computing, and not much about wireless communication. *Private 5G: A Systems Approach*—this book—changed that. I learned what I needed to know about the radio access network, and how the pieces come together to be more than the sum of their parts. More than that, the book shows that, despite the complexity of the 3GPP cellular standards and the long and frustrating history of closed network equipment in the cellular networking space, open source software is gaining a foothold. Open source platforms like Aether and Magma—both discussed in this book—are seeing practical deployment in a wide variety of settings. The book even has a guide for readers to bring up the Aether platform, including a 5G small cell.

For me, then, the story comes full circle. Private 5G networks are something you can touch, code, and deploy yourself. Armed with the knowledge of how 5G access networks work, and with hands-on experience with open source software, just imagine the places you'll go!

Jennifer Rexford
Princeton, New Jersey

Preface

When we wrote our introductory 5G book three years ago, our goal was to help people with experience building Internet and cloud services to understand the opportunity to bring best practices from those systems to the mobile cellular network. On paper (and in the press) 5G had set an ambitious goal of transformative changes, adopting a cloud-inspired architecture and supporting a new set of innovative services. But the gap between that aspirational story and the reality of 40 years of network operators and hardware vendors protecting their incumbent advantages made for a challenging pivot. So we started with the basics, and set out to explain the fundamental networking concepts and design principles behind the myriad of acronyms that dominate mobile cellular networking.

Because 5G adopts many of the principles of cloud native systems, it promises to bring the feature velocity of the cloud to Telco environments. That promise is being delivered most successfully in private 5G deployments that are less constrained by existing Telco organizations and legacy infrastructure. What started out as sketches on a whiteboard three years ago is now becoming a reality: Several cloud providers are offering private 5G solutions for enterprises, and there is a complete open source implementation of a 5G-enabled edge cloud that the Internet community can learn from and build upon.

The architecture described in this book is not limited to private deployments. It includes the necessary background information about the mobile cellular network, much of which is rooted in its origin story as a Telco voice network, but the overarching theme is to describe the network through the lens of private deployments of 5G connectivity as a managed cloud service. This includes adopting best practices in horizontally scalable microservices, Software-

Defined Networking (SDN), and cloud operational practices such as DevOps. These practices are appropriate for traditional operators, cloud providers, and enterprises alike, but it is emerging use cases in private deployments that will benefit first.

The book makes extensive use of open source software—specifically, the Aether and Magma projects—to illustrate how Private 5G can be realized in practice. The availability of open software informs our understanding of what has historically been a proprietary and opaque system. The result complements the low-level engineering documents that are available online (and to which we provide links) with an architectural roadmap for anyone trying to understand all the moving parts, how they fit together, and how they can be operationalized. And once you're done reading the book, we encourage you to jump into the hands-on appendix that walks you through the step-by-step process of deploying that software in your own local computing environment.

Acknowledgements


The software described in this book is due to the hard work of the ONF engineering team, the Magma engineering team, and the open source communities that work with them. Bilal Saleem did the heavy lifting on Aether OnRamp (described in the *About the Software* Appendix), with a special thanks to Ajay Thakur, Andy Bavier, Gabriel Arrobo, and Muhammad Shahbaz for their guidance and feedback.

Thanks to the members of the community who contributed text or corrections to the book, including:

- Edmar Candeia Gurjão
- Muhammad Shahbaz
- Mugahed Izzeldin
- Robert MacDavid
- Simon Leinen
- Tiago Barros
- Gabriel Arrobo
- Ajay Thakur

The picture of a Magma deployment in Chapter 5 was provided by Shaddi Hasan.

Larry Peterson, Oguz Sunay, and Bruce Davie
April 2023



Chapter 1: Introduction

Mobile networks, which have a 40-year history that parallels the Internet's, have undergone significant change. The first two generations supported voice and then text, with 3G defining the

transition to broadband access, supporting data rates measured in hundreds of kilobits per second. Today, the industry is transitioning from 4G (with data rates typically measured in the few megabits per second) to 5G, with the promise of a tenfold increase in data rates.

But 5G is about much more than increased bandwidth. 5G represents a fundamental rearchitecting of the access network in a way that leverages several key technology trends and sets it on a path to enable much greater innovation. In the same way that 3G defined the transition from voice to broadband, 5G's promise is primarily about the transition from a single access service (broadband connectivity) to a richer collection of edge services and devices. 5G is expected to provide support for immersive user interfaces (e.g., Augmented Reality, Virtual Reality), mission-critical applications (e.g., public safety, autonomous vehicles), and the Internet of Things (IoT). Because these use cases will include everything from home appliances to industrial robots to self-driving cars, 5G will support not only humans accessing the Internet from their smartphones, but also swarms of autonomous devices working together on their behalf.

There is more to supporting these services than just improving bandwidth or latency to individual users. As we will see, a fundamentally different edge network architecture is required. The requirements for this architecture are ambitious, and can be illustrated by three classes of capabilities:

- To support *Massive Internet of Things*, potentially including devices with ultra-low energy (10+ years of battery life), ultra-low complexity (10s of bits per second), and ultra-high density (1 million nodes per square kilometer).
- To support *Mission-Critical Control*, potentially including ultra-high availability (greater than 99.999% or “five nines”), ultra-low latency (as low as 1 ms), and extreme mobility (up to 100 km/h).
- To support *Enhanced Mobile Broadband*, potentially including extreme data rates (multi-Gbps peak, 100+ Mbps sustained) and extreme capacity (10 Tbps of aggregate throughput per square kilometer).

These targets will certainly not be met overnight, but that's in keeping with each generation of the mobile network being a decade-long endeavor.

On top of these quantitative improvements to the capabilities of the access network, 5G is being viewed as a chance for building a platform to support innovation. Whereas prior access networks were generally optimized for known services (such as voice calls and SMS), the Internet has been hugely successful in large part because it supported a wide range of applications that were not even thought of when it was first designed. The 5G network is designed with this same goal: enabling future applications beyond those we fully recognize today. For an example of the grand vision for 5G, see the whitepaper from one of the industry leaders.

Further Reading

Qualcomm Whitepaper. [Making 5G NR a Reality](#). December 2016.

The 5G mobile network, because it is on an evolutionary path and not a point solution, includes standardized specifications, a range of implementation choices, and a long list of aspirational goals. Because this leaves so much room for interpretation, our approach to describing 5G is grounded in three mutually supportive principles. The first is to apply a *systems lens*, which is to say, we explain the sequence of design decisions that lead to a solution rather than fall back on enumerating the overwhelming number of acronyms or individual point technologies as a *fait accompli*. The second is to aggressively disaggregate the system. Building a disaggregated, virtualized, and software-defined 5G access network is the direction the industry is already headed (for good technical and business reasons), but breaking the 5G network down into its elemental components is also the best way to explain how 5G works.

The third principle is to illustrate how 5G can be realized in practice by drawing on specific engineering decisions made in an open source implementation. This implementation leverages best practices in building cloud apps, which is an essential aspect of 5G evolving into a platform for new services. This implementation also targets enterprises that are increasingly deploying 5G locally, and

using it to help automate their manufacturing, retail, and business practices—a trend that has been dubbed *Industry 4.0*. Such enterprise-level deployments are known as *Private 5G*, but there is nothing about the technical approach that couldn't be adopted throughout the more traditional “public mobile network” that comes to mind when you think about your cell service today. The only difference is that private deployments are more aggressively embracing the cloud practices that will ultimately distinguish 5G from earlier generations.

Further Reading

K. Schwab. [The Fourth Industrial Revolution](#). World Economic Forum.

What this all means is that there is no simple definition of 5G, any more than there is for the Internet. It is a complex and evolving system, constrained by a set of standards that purposely give all the stakeholders many degrees of freedom. In the chapters that follow, it should be clear from the context whether we are talking about *standards* (what everyone must do to interoperate), *trends* (where the industry seems to be headed), or *implementation choices* (examples to make the discussion more concrete). By adopting a systems perspective throughout, our intent is to describe 5G in a way that helps the reader navigate this rich and rapidly evolving system.

1.1 Standardization Landscape

As of 3G, the generational designation corresponds to a standard defined by the *3rd Generation Partnership Project (3GPP)*. Even though its name has “3G” in it, the 3GPP continues to define the standards for 4G, 5G, and so on, each of which corresponds to a sequence of releases of the standard. Release 15 is considered the demarcation point between 4G and 5G, with Release 17 having been completed in 2022.

In addition to 3GPP-defined standards, national governments establish how the radio spectrum is used locally. Unlike Wi-Fi, for which there is international agreement that permits anyone to use a channel at either 2.4 or 5 GHz (these are unlicensed bands), governments have auctioned off and licensed exclusive use of various frequency bands to service providers, who in turn sell mobile access service to their subscribers. The use of licensed spectrum brings certain benefits such as greater control over the quality of service delivered, while also imposing costs both in terms of paying for licenses and in the complexity of the systems needed to manage access to the spectrum. We will explore how these costs and benefits play out in subsequent chapters.

There is also a shared-license band at 3.5 GHz, called *Citizens Broadband Radio Service (CBRS)*, set aside in North America for cellular use. Similar spectrum is being set aside in other countries. The CBRS band allows 3 tiers of users to share the spectrum: first right of use goes to the original owners of this spectrum (naval radars and satellite ground stations); followed by priority users who receive this right over 10MHz bands for three years via regional auctions; and finally the rest of the population, who can access and utilize a portion of this band as long as they first check with a central database of registered users. CBRS, along with standardization efforts to extend mobile cellular networks to operate in the unlicensed bands, opens the door for private cellular networks similar to Wi-Fi. This is proving especially attractive to enterprises looking to establish a *Private 5G* service.

The specific frequency bands that are licensed for cellular networks vary around the world, and are complicated by the fact that network operators often simultaneously support both old/legacy technologies and new/next-generation technologies, each of which occupies a different frequency band. The high-level summary is that traditional cellular technologies range from 700-2400 MHz, with new mid-spectrum allocations now happening at 6 GHz, and millimeter-wave (mmWave) allocations opening above 24 GHz.

While the specific frequency band is not directly relevant to understanding 5G from an architectural perspective, it does impact the physical-layer components, which in turn has indirect

ramifications on the overall 5G system. We identify and explain these ramifications in later chapters, keeping in mind that ensuring the allocated spectrum is used *efficiently* is a critical design goal.

Finally, in addition to the long-established 3GPP standards body and the set of national regulatory agencies around the world, a new organization—called the *Open-RAN Alliance (O-RAN)*—has recently been established to focus on “opening up the Radio Access Network”. We’ll see specifically what this means and how the O-RAN differs from the 3GPP in Chapter 4, but for now, its existence highlights an important dynamic in the industry: 3GPP has become a vendor-dominated organization, with network operators (AT&T and China Mobile were the founding members) creating O-RAN to break vendor lock-in.

1.2 Access Networks

The mobile cellular network is part of the access network that implements the Internet’s so-called *last mile*. (Another common access technology is *Passive Optical Networks (PON)*, colloquially known as Fiber-to-the-Home.) These mobile access networks have historically been provided by both big and small *Mobile Network Operators (MNOs)*. Global MNOs such as AT&T run access networks at thousands of aggregation points of presence across a country such as the US, along with a national backbone that interconnects those sites. Small regional and municipal MNOs might run an access network with one or two points of presence, and then connect to the rest of the Internet through some large operator’s backbone.

Figure 1. A global mobile network built by first aggregating traffic from hundreds of wireless base stations, and then interconnecting those aggregation points over the Internet.

As illustrated in [Figure 1](#), access networks are physically anchored at thousands of aggregation points of presence within close proximity to end users, each of which serves anywhere from 1,000-100,000 subscribers, depending on population density. In practice, the physical deployment of these “edge” locations vary from operator to operator, but one possible scenario is to anchor

both the cellular and wireline access networks in Telco *Central Offices*.

Historically, the Central Office—officially known as the *PSTN (Public Switched Telephone Network) Central Office*—anchored wired access (both telephony and broadband), while the cellular network evolved independently by deploying a parallel set of *Mobile Telephone Switching Offices (MTSO)*. Each MTSO serves as a *mobile aggregation* point for the set of cell towers in a given geographic area. For our purposes, the important idea is that such aggregation points exist, and it is reasonable to think of them as defining the edge of an operator-managed access network. For simplicity, we sometimes use the term “Central Office” as a synonym for both types of edge sites.

Finally, one aspect of the mobile network that may not be obvious from [Figure 1](#) is that it supports global connectivity, independent of the Internet (which is technically just one of many available backbone technologies). That is, the cellular network supports a universal addressing scheme, similar in principle (but significantly different in details) from the Internet’s universal IP-based addressing scheme. This addressing scheme makes it possible to establish a voice call between any two cell phones, but of course, IP addresses still come into play when trying to establish a data (broadband) connection to/from a cell phone or other mobile device. Understanding the relationship between mobile addresses and IP addresses is a topic we will explore in later chapters.

1.3 Managed Cloud Service

The previous section gives a decidedly Telco-centric view of the mobile cellular network, which makes sense because Telcos have been the dominant MNOs for the past 40+ years. But with 5G’s focus on broadening the set of services it supports, and embracing general platforms that can host yet-to-be-invented applications, the mobile cellular network is starting to blur the line between the access network and the cloud.

5G, Wi-Fi, and the Role of Spectrum

WiFi networks use unlicensed radio spectrum that do not require WiFi network operators to get advance regulatory approval. At the same time, anyone can access the same spectrum, subject to limits on transmission power. As a result, WiFi networks share their bands with devices including baby monitors, cordless phones, etc., so the WiFi MAC layer assumes the presence of physical-layer interference. Enterprise WiFi deployments, such as those on college campuses and in corporate office buildings, perform more centralized management of interference across multiple overlapping access points, but risk of interference remains and thus the service remains best-effort.

Cellular access networks typically use licensed spectrum that is owned or leased by the carrier for long periods of time at high cost. Even “lightly licensed” spectrum such as CBRS offers more control over interference than Wi-Fi. Since the cellular radio has exclusive access to spectrum over a geographic region, cellular waveforms are designed for wide-area coverage and high spectral efficiency. Managing access to the spectrum, as we shall see, is an important aspect of the 5G architecture.

Many of the differences between 5G and Wi-Fi follow from the differences in spectrum and radio characteristics. For example, cellular deployments, with the expense of spectrum being a given, have historically been carried out by well-resourced actors who can acquire land, build and connect towers, and hire skilled staff. However, the rise of enterprise 5G and the availability of lightly licensed spectrum is leading to a blurring of the lines between the two approaches.

The rest of this book explains what that means in detail. As an overview, thinking of 5G connectivity as a cloud service means that instead of using purpose-built devices and telephony-based operational practices to deliver mobile connectivity, the 5G network is built from commodity hardware, software-defined networks, and cloud-based operational practices. And, just as with familiar cloud applications, the end result is a system that increases both feature velocity and operational uniformity. These advantages are available to legacy MNOs, but whether they will fully embrace them is yet to be seen, so we do not limit ourselves to existing stakeholders or business models. In particular, this book focuses on how enterprises can be their own MNOs, or alternatively, acquire private 5G connectivity as a managed cloud service from non-traditional

MNOs.

To this end, [Figure 2](#) depicts a simplified Private 5G deployment that the rest of this book works toward. At a high level, the figure shows a wide range of enterprise use cases that might take advantage of 5G connectivity, with the data plane of the 5G service running on-prem (on an edge cloud running within the enterprise), and the control plane of the 5G service running off-prem (in the global cloud).[\[1\]](#) Enterprise administrators control their service through a management console, much in the same way they might log into an AWS, GCP, or Azure console to control a cloud-based storage or compute service. Finally, applications are distributed across both edge and centralized clouds, taking advantage of what is commonly referred to as a *hybrid cloud*.

We use the term “data plane” in the generic sense in this description. As we’ll see in Chapter 2, the 5G architecture refers to it as “user plane”.

Figure 2. Enterprise-based deployment of 5G connectivity, running as a managed cloud service.

Hosting a 5G connectivity service on an edge cloud is perfectly aligned with one of the most pronounced trends in cloud computing: moving elements of the cloud from the datacenter to locations that are in close proximity to end users and their devices. Before looking at how to realize 5G on an edge cloud, we start by considering why edge clouds are gaining momentum in the first place.

The cloud began as a collection of warehouse-sized datacenters, each of which provided a cost-effective way to power, cool, and operate a scalable number of servers. Over time, this shared infrastructure lowered the barrier to deploying scalable Internet services, but today, there is increasing pressure to offer low-latency/high-bandwidth cloud applications that cannot be effectively implemented in remote datacenters. Augmented Reality (AR), Virtual Reality (VR), Internet of Things (IoT), and Autonomous Vehicles are all examples of this kind of application. Such applications benefit from moving at least part of their functionality out of the datacenter and towards the edge of the

network, closer to end users.

The idea of such deployments is to first collect operational data on assets and infrastructure, from sensors, video feeds and telemetry from machinery. It then applies Machine Learning (ML) or other forms of analysis to this data to gain insights, identify patterns and predict outcomes (e.g., when a device is likely to fail). The final step is to automate industrial processes so as to minimize human intervention and enable remote operations (e.g., power optimization, idling quiescent machinery). The overall goal is to create an IT foundation for continually improving industrial operations through software.

But precisely where this edge is *physically* located depends on who you ask. If you ask a network operator that already owns and operates thousands of Central Offices, then their Central Offices are an obvious answer. Others might claim the edge is located at the 14,000 Starbucks locations (for example) across the US, and still others might point to the tens of thousands of cell towers spread across the globe. Our approach is to be location agnostic, but to make the discussion concrete, we use enterprises as our exemplar deployment.

At the same time cloud providers started pursuing edge deployments, network operators began to re-architect their access network to use the same commodity hardware and best practices in building scalable software as the cloud providers. Such a design, which is sometimes referred to as CORD (*Central Office Re-architected as a Datacenter*), supports both the access network and edge services co-located on a shared cloud platform. This platform is then replicated across hundreds or thousands of operator sites, including Central Offices.

Traditional network operators did this because they wanted to take advantage of the same economies of scale and feature velocity as cloud providers. CORD gave them a general architecture to work towards, but also an open source Kubernetes-based reference implementation to model their solutions on. That original implementation of CORD is the direct predecessor to the Aether platform that we use as a reference implementation in this book.

Further Reading

L. Peterson, *et al.* [Central Office Re-architected as a Datacenter.](#).. IEEE Communications, October 2016.

A.D. Little Report. [Who Dares Wins! How Access Transformation Can Fast-Track Evolution of Operator Production Platforms.](#) September 2019.

An important takeaway from this discussion is that to understand how 5G is being implemented, it is helpful to have a working understanding of how clouds are built. This includes the use of *commodity hardware* (both servers and bare-metal switches), horizontally scalable *microservices* (also referred to as *cloud native*), and *Software-Defined Networks (SDN)*. It is also helpful to have an appreciation for how cloud software is developed, tested, deployed, and operated, including practices such as *DevOps* and *Continuous Integration / Continuous Deployment (CI/CD)*. We recommend two companion books to help fill the gaps in your understanding of these foundational technologies.

Further Reading

[Software-Defined Networks: A Systems Approach.](#) November 2021.

[Edge Cloud Operations: A Systems Approach.](#) June 2022.

1.4 Beyond 5G

From the moment MNOs started rolling out 5G in 2019, people started talking about what comes next. The obvious answer is 6G, but it's not at all clear that the decadal generations of the past 40 years will continue into the future. Today, you hear alternatives like "NextG" and "Beyond 5G" more often than 6G, which could be a sign that the industry is undergoing a fundamental shift. And there is an argument that we're in the midst of a sea change that will render the generational distinction largely meaningless. There are two complementary reasons for this, both at the heart of what's important about Private 5G.




The first factor is that by adopting cloud technologies, the mobile cellular network is hoping to cash in on the promise of feature velocity. This “agility” story was always included in the early 5G promotional material, as part of the case for why a 5G upgrade would be a worthwhile investment, but the consequence of those technologies now finding their way into the mainstream is that new features can be introduced rapidly and deployed continuously. At some point, the frequency of continual improvements renders generational distinctions irrelevant.




The second factor is that agility isn’t only about cadence; it’s also about customization. That is, these changes can be introduced bottom-up—for example by enterprises and their edge cloud partners in the case of Private 5G—without necessarily depending on (or waiting for) a global standardization effort. If an enterprise finds a new use case that requires a specialized deployment, only its Private 5G deployment needs to adopt the necessary changes. Reaching agreement with all the incumbent stakeholders will no longer be a requirement.




It’s anyone’s guess where this will take us, but it will be interesting to see how this dynamic impacts the role of standardization: what aspects of the mobile network require global agreement and what aspects do not because they can evolve on a case-by-case basis. While standards often spur innovation (TCP and HTTP are two great examples from the Internet experience), sometimes standards actually serve as a barrier to competition, and hence, innovation. Now that software is eating the mobile cellular network—with Private 5G deployed in enterprises likely setting the pace—we will learn which standards are which.



In summary, that 5G is on an evolutionary path is the central theme of this book. We call attention to its importance here, and revisit the topic throughout the book. We are writing this book for *system generalists*, with the goal of helping bring a community that understands a broad range of systems issues (but knows little or nothing about the cellular network) up to speed so they can play a role in its evolution. This is a community that understands both feature velocity and best practices in building robust scalable systems, and so has an important role to play in bringing all of 5G’s


potential to fruition.











Chapter 2: Architecture

This chapter identifies the main architectural components of the

mobile cellular network. We need to introduce some terminology to do this, which can be confusing for those whose networking background comes from the Internet. This is partly because some of what needs to happen in a mobile network, such as keeping track of which base station is serving a given mobile device, doesn't really have a parallel in fixed networks. On top of that, the terminology came out of the 3GPP standardization process, which was historically concerned with telephony and almost completely disconnected from the IETF and other Internet-related efforts. To further confuse matters, 3GPP terminology often changes with each generation (e.g., a base station is called eNB in 4G and gNB in 5G). We address situations like this by using generic terminology (e.g., base station), and referencing the 3GPP-specific counterpart only when the distinction is helpful. This example is only the tip of the terminology iceberg. Marcin Dryjanski's blog post gives a broader perspective on the complexity of terminology in 5G.

Further Reading

Marcin Dryjanski. [LTE and 5G Differences: System Complexity](#). July 2018.

2.1 Overview

The mobile cellular network provides wireless connectivity to devices that are (potentially) on the move. These devices, which are known as *User Equipment (UE)*, have traditionally corresponded to mobile phones and tablets, but increasingly include cars, drones, industrial and agricultural machines, robots, home appliances, medical devices, and so on. In some cases, the UEs may be devices that do not move, e.g., router interfaces using cellular connectivity to provide broadband access to remote dwellings.

Figure 3. Mobile cellular networks consist of a Radio Access Network (RAN) and a Mobile Core.

As shown in [Figure 3](#), the mobile cellular network consists of two main subsystems: the *Radio Access Network (RAN)* and the *Mobile Core*. The RAN manages the radio resources (i.e., spectrum), making

sure it is used efficiently and meets the quality of service (QoS) requirements of every user. It corresponds to a distributed collection of base stations. As noted above, these are cryptically named *eNodeB* or *eNB* (which is short for *evolved Node B*) in 4G. In 5G, base stations are known as *gNB*, where the “g” stands for *next Generation*.

The Mobile Core is a bundle of functionality (conventionally packaged as one or more devices) that serves several purposes.

- Authenticates devices prior to attaching them to the network
- Provides Internet (IP) connectivity for both data and voice services.
- Ensures this connectivity fulfills the promised QoS requirements.
- Tracks user mobility to ensure uninterrupted service.
- Tracks subscriber usage for billing and charging.

For readers familiar with the Internet architecture and Wi-Fi as a common access technology, some of these functions might look a bit surprising. For example, Wi-Fi, like most of the Internet, normally provides a best-effort service, whereas cellular networks often aim to deliver some sort of QoS guarantee. Tracking subscribers for both mobility and billing are also not the sort of things we tend to think about in the Internet, but they are considered important functions for cellular networks. The reasons for these differences are numerous, including the typically large costs of acquiring cellular spectrum and maintaining the infrastructure to use it such as radio towers. With that large investment, there is a desire to recoup costs by charging subscribers, which in turn leads to making some sort of service guarantees to those subscribers to justify the cost. There is also a need to maximize the efficiency of spectrum usage. Much of the complexity of the mobile core follows from these requirements being imposed by service providers. Even when we get to enterprises running their own 5G networks, they still need to manage the usage of spectrum to obtain the benefits of 5G over Wi-Fi, such as more predictable control over latency and bandwidth.

Note that Mobile Core is another example of a generic term. In 4G it was called the *Evolved Packet Core (EPC)* and in 5G it is called the

5G Core (5GC). Moreover, even though the word “Core” is in its name, the Mobile Core runs near the edge of the network, effectively providing a bridge between the RAN in some geographic area and the greater IP-based Internet. 3GPP provides significant flexibility in how the Mobile Core is geographically deployed, ranging from minimal deployments (the RAN and the mobile core can be co-located) to areas that are hundreds of kilometers wide. A common model is that an instantiation of the Mobile Core serves a metropolitan area. The corresponding RAN would then span several dozens (or even hundreds) of cell towers in that geographic area.

Taking a closer look at [Figure 3](#), we see that a *Backhaul Network* interconnects the base stations that implement the RAN with the Mobile Core. This network is typically wired, may or may not have the ring topology shown in the figure, and is often constructed from commodity components found elsewhere in the Internet. For example, the *Passive Optical Network (PON)* that implements Fiber-to-the-Home is a prime candidate for implementing the RAN backhaul, with the RAN effectively running as an *overlay* on top of whatever technology is used. Switched ethernet, such as you might find in an enterprise, is another suitable choice. The backhaul network is obviously a necessary part of the RAN, but it is an implementation choice and not prescribed by the 3GPP standard.

Although 3GPP specifies all the elements that implement the RAN and Mobile Core in an open standard—including sub-layers we have not yet introduced—network operators have historically bought proprietary implementations of each subsystem from a single vendor. This lack of an open source implementation contributes to the perceived “opaqueness” of the mobile cellular network in general, and the RAN in particular. And while it is true that base stations contain sophisticated algorithms for scheduling transmission on the radio spectrum—algorithms that are considered valuable intellectual property of the equipment vendors—there is significant opportunity to open and disaggregate both the RAN and the Mobile Core. This book gives a recipe for how to do exactly that.

Before getting to those details, we have three more architectural concepts to introduce. First, [Figure 4](#) redraws components from

Figure 3 to highlight the fact that a base station has an analog component (depicted by an antenna) and a digital component (depicted by a processor pair). This book mostly focuses on the latter, but we introduce enough information about the over-the-air radio transmission to appreciate its impact on the overall architecture.

Figure 4. Mobile Core divided into a Control Plane and a User Plane, an architectural feature known as CUPS: Control and User Plane Separation.

The second concept, also depicted in **Figure 4**, is to partition the Mobile Core into a *Control Plane* and *User Plane*. This is similar to the control/data plane split that anyone familiar with the Internet would recognize, and draws in particular on the ideas of software-defined networking (SDN) by placing control and user planes in separate devices. 3GPP has introduced a corresponding acronym—*CUPS, Control and User Plane Separation*—to denote this idea. One motivation for CUPS is to enable control plane resources and data plane resources to be scaled independently of each other.

Finally, one of the key aspirational goals of 5G is the ability to segregate traffic for different usage domains into isolated *network slices*, each of which delivers a different level of service to a collection of devices and applications. Thinking of a network slice as a wireless version of a virtual network is a fair approximation, although as we'll see in later chapters, the implementation details differ.

Figure 5. Different usage domains (e.g., IoT and Video Streaming) instantiate distinct network slices to connect a set of devices with one or more applications.

For example, **Figure 5** shows two slices, one supporting IoT workloads and the other supporting multimedia streaming traffic. As we'll see throughout the book, an important question is how slicing is realized end-to-end, across the radio, the RAN, and the Mobile Core. This is done through a combination of allocating distinct resources to each slice and scheduling shared resources on behalf of a set of slices.

2.2 Radio Transmission

Before describing the RAN and Mobile Core subsystems, we first call attention to the obvious: that the base stations that comprise the RAN communicate with UEs via electromagnetic radio waves. This book is not about the physics of this over-the-air communication, and only skims the surface of the information theory that underlies it. But identifying the abstract properties of wireless communication is an essential foundation for understanding the rest of the 5G architecture.

If you imagine the base stations as implementing a multi-layer protocol stack (which, as we'll see in Chapter 4, they do), then radio transmission is the responsibility of the bottom-most layers of that stack, where (a) digital/analog conversion happens, and (b) analog radio waves are transmitted/received. Chapter 3 introduces radio transmission with enough specificity to lay the necessary foundation, so we're able to understand all the layers that come above it.

For the purposes of this chapter, we only need to understand the following. First, the RAN is responsible for managing how the radio spectrum is shared among thousands of UEs connected to hundreds of base stations in a geographic region. The primary purpose of Chapter 3 is to establish an abstract interface by which the RAN can manage that spectrum without having to worry about the details of waveforms, modulation, or coding algorithms. All important topics, to be sure, but in the realm of information theory rather than system design that is the focus of this book.

Figure 6. Abstractly, measures of signal quality (CQI) and declarations of intended data delivery quality (5QI) are passed up and down the RAN stack.

Second, there are two important pieces of information shared between the higher layers of the base station protocol stack that manages the RAN as a whole, and the lower layers of the stack that manage radio transmissions on a particular base station. One is the signal-to-noise ratio that the base station observes when communicating with each UE. This is called the *Channel Quality*

Indicator (CQI) and it is passed *up* from the radio. The other is the quality of service the network wants to give a particular UE. This is called the *5G QoS Identifier (5QI)* and it is passed *down* to the radio. This abstract summary, as shown in [Figure 6](#), is sufficient to introduce the RAN and Mobile Core. We will fill in more details about both of these parameters in Chapter 3.

Uniqueness of Wireless Links

While it is common in networking to abstract the link layer by treating the link as something that just delivers packets at some rate from point A to point B, there are important differences between wireless links and wired links that cannot be entirely abstracted away at higher layers. This is especially true when mobile devices are involved, as the quality of a link will vary depending on the distance between transmitter and receiver, the relative velocity of the endpoints, reflections of radio waves from other objects, and interference from other transmitters. All of these factors come into play in determining the Channel Quality Indicator (CQI).

Further complicating the picture in a mobile network is that a given UE is often within reach of more than one base station, presenting the option to handover the UE from one base station to another. The decision to do so is not just a matter of picking the base station with the best channel quality, but rather a matter of trying to optimize the whole system, in which the goal is to support as many UEs as possible at the desired quality level given the available spectrum and coverage.

Finally, like the rest of the mobile cellular network, the radio comes with a set of acronyms, with *LTE (Long-Term Evolution)* and *NR (New Radio)* being the two most widely known. These are marketing terms commonly associated with the radio technology for 4G and 5G, respectively. They are important only in the sense that many of the new features promised by 5G can be directly attributed to improvements in the underlying radio technology. For our purposes, the key is the set of new *use cases* the upgraded radio technology enables, and why. We introduce these improvements to the radio in Chapter 3, and tie them to the use cases they enable. Subsequent chapters will then explain how the RAN and Mobile Core need to evolve so as to deliver on this potential.

2.3 Radio Access Network

We now describe the RAN by sketching the role each base station plays. Keep in mind this is like describing the Internet by explaining how a router works—not an unreasonable place to start, but it doesn’t fully do justice to the end-to-end story.

First, each base station establishes the wireless channel for a subscriber’s UE upon power-up or upon handover when the UE is active. This channel is released when the UE remains idle for a predetermined period of time. Using 3GPP terminology, this wireless channel is said to provide a *radio bearer*. The term “bearer” has historically been used in telecommunications (including early wireline technologies such as ISDN) to denote a data channel, as opposed to a channel that carries signaling information.

Figure 7. UE detects (and connects to) base station.

Second, each base station establishes “3GPP Control Plane” connectivity between the UE and the corresponding Mobile Core Control Plane component, and forwards signaling traffic between the two. This signaling traffic enables UE authentication, registration, and mobility tracking.

Figure 8. Base Station establishes control plane connectivity between each UE and the Mobile Core.

Third, for each active UE, the base station establishes one or more tunnels to the corresponding Mobile Core User Plane component. [Figure 9](#) shows just two (one for voice and one for data), and while in practice 4G was limited to just two, 5G aspires to support many such tunnels as part of a generalized network slicing mechanism.

Figure 9. Base station establishes one or more tunnels between each UE and the Mobile Core’s User Plane (known in 3GPP terms as PDU session).

Fourth, the base station forwards both control and user plane packets between the Mobile Core and the UE. These packets are tunneled over SCTP/IP and GTP/UDP/IP, respectively. SCTP

(Stream Control Transport Protocol) is an alternative reliable transport to TCP, tailored to carry signaling (control) information for telephony services. GTP (a nested acronym corresponding to (General Packet Radio Service) Tunneling Protocol) is a 3GPP-specific tunneling protocol designed to run over UDP.

It is noteworthy that connectivity between the RAN and the Mobile Core is IP-based. This was introduced as one of the main changes between 3G and 4G. Prior to 4G, the internals of the cellular network were circuit-based, which is not surprising given its origins as a voice network. This also helps to explain why in Section 2.1 we characterized the RAN Backhaul as an overlay running on top of some Layer 2 technology.

Figure 10. Base Station to Mobile Core (and Base Station to Base Station) control plane tunneled over SCTP/IP and user plane tunneled over GTP/UDP/IP.

Fifth, each base station coordinates UE handovers with neighboring base stations, using direct station-to-station links. Exactly like the station-to-core connectivity shown in the previous figure, these links are used to transfer both control plane (SCTP over IP) and user plane (GTP over UDP/IP) packets. The decision as to when to do a handover is based on the CQI values being reported by the radio on each of the base stations within range of the UE, coupled with the 5QI value those base stations know the RAN has promised to deliver to the UE.

Figure 11. Base Stations cooperate to implement UE hand over.

Sixth, the base stations coordinate wireless multi-point transmission to a UE from multiple base stations, which may or may not be part of a UE handover from one base station to another.

Figure 12. Base Stations cooperate to implement multipath transmission (link aggregation) to UEs.

The main takeaway is that the base station can be viewed as a specialized forwarder. In the Internet-to-UE direction, it fragments outgoing IP packets into physical layer segments and schedules them for transmission over the available radio spectrum, and in the

UE-to-Internet direction it assembles physical layer segments into IP packets and forwards them (over a GTP/UDP/IP tunnel) to the upstream user plane of the Mobile Core. Also, based on observations of the wireless channel quality and per-subscriber policies, it decides whether to (a) forward outgoing packets directly to the UE, (b) indirectly forward packets to the UE via a neighboring base station, or (c) utilize multiple paths to reach the UE. The third case has the option of either spreading the physical payloads across multiple base stations or across multiple carrier frequencies of a single base station (including Wi-Fi).

In other words, the RAN as a whole (i.e., not just a single base station) not only supports handovers (an obvious requirement for mobility), but also *link aggregation* and *load balancing*, mechanisms that are similar to those found in other types of networks. These functions imply a global decision-making process, whereby it's possible to forward traffic to a different base station (or to multiple base stations) in an effort to make efficient use of the radio spectrum over a larger geographic area. We will revisit how such RAN-wide (global) decisions can be made using SDN techniques in Chapter 4.

2.4 Mobile Core

At the most basic level, the function of the Mobile Core is to provide packet data network connectivity to mobile subscribers, i.e., connect them to the Internet. (The mobile network assumes that multiple packet data networks can exist, but in practice the Internet is the one that matters). As we noted above, there is more to providing this connectivity than meets the eye: the Mobile Core ensures that subscribers are authenticated and aims to deliver the service qualities to which they have subscribed. As subscribers may move around among base station coverage areas, the Mobile Core needs to keep track of their whereabouts at the granularity of the serving base station. The reasons for this tracking are discussed further in Chapter 5. It is this support for security, mobility, and QoS that differentiates the cellular network from Wi-Fi.

We start with the security architecture, which is grounded in two

trust assumptions. First, each base station trusts that it is connected to the Mobile Core by a secure private network, over which it establishes the tunnels introduced in [Figure 10](#): a GTP/UDP/IP tunnel to the Core's User Plane (Core-UP) and a SCTP/IP tunnel to the Core's Control Plane (Core-CP). Second, each UE has an operator-provided SIM card, which contains information that uniquely identifies the subscriber and includes a secret key that the UE uses to authenticate itself.

The identifier burned into each SIM card, called an *IMSI* (*International Mobile Subscriber Identity*), is a globally unique identifier for every device connected to the global mobile network. Each IMSI is a 64-bit, self-describing identifier, which is to say, it includes a *Format* field that effectively serves as a mask for extracting other relevant fields. For example, the following is the interpretation we assume in this book (where IMSIs are commonly represented as an up to 15-digit decimal number):

- **MCC**: Mobile Country Code (3-digit decimal number).
- **MNC**: Mobile Network Code (2 or 3-digit decimal number).
- **ENT**: Enterprise Code (3-digit decimal number).
- **SUB**: Subscriber (6-digit decimal number).

The first two fields (*MCC*, *MNC*) are universally understood to uniquely identify the MNO, while that last two fields are one example of how an MNO might use additional hierarchical structure to uniquely identify every device it serves. We are working towards delivering 5G connectivity to enterprises (hence the *ENT* field), but other MNOs might assign the last 9 digits using some other structure.

The *MCC/MNC* pair—which is also called the *Public Land Mobile Network (PLMN)* identifier—plays a role in roaming: when a UE tries to connect to a “foreign network” those fields are used to find the “home network”, where the rest of the IMSI leads to a subscriber profile that says whether or not roaming is enabled for this device. The following walks through what happens when a device connects to its home network; more information about the global ramifications is given at the end of the section.

Figure 13. Sequence of steps to establish secure Control and User

Plane channels.

Figure 13 shows the per-UE connection sequence. When a UE first becomes active, it communicates with a nearby base station over a temporary (unauthenticated) radio link (Step 1). The base station forwards the request to the Core-CP over the existing SCTP connection, and the Core-CP (assuming it recognizes the IMSI) initiates an authentication protocol with the UE (Step 2). 3GPP identifies a set of options for authentication and encryption, where the actual protocols used are an implementation choice. For example, *Advanced Encryption Standard* (AES) is one of the options for encryption. Note that this authentication exchange is initially in the clear since the base station to UE link is not yet secure.

Once the UE and Core-CP are satisfied with each other's identity, the Core-CP informs the other 5GC components of the parameters they will need to service the UE (Step 3). This includes: (a) instructing the Core-UP to initialize the user plane (e.g., assign an IP address to the UE and set the appropriate 5QI); (b) instructing the base station to establish an encrypted channel to the UE; and (c) giving the UE the symmetric key it will need to use the encrypted channel with the base station. The symmetric key is encrypted using the public key of the UE (so only the UE can decrypt it, using its secret key). Once complete, the UE can use the end-to-end user plane channel through the Core-UP (Step 4).

There are three additional details of note about this process. First, the secure control channel between the UE and the Core-CP set up during Step 2 remains available, and is used by the Core-CP to send additional control instructions to the UE during the course of the session. In other words, unlike the Internet, the network is able to “control” the communication settings in edge devices.

Second, the user plane channel established during Step 4 is referred to as the *Default Data Radio Bearer*, but additional channels can be established between the UE and Core-UP, each with a potentially different 5QI. This might be done on an application-by-application basis, for example, based on policies present in the Core-CP for packets that require special/different treatment.

Figure 14. Sequence of per-hop tunnels involved in an end-to-end

User Plane channel.

In practice, these per-flow tunnels are often bundled into a single inter-component tunnel, which makes it impossible to differentiate the level of service given to any particular end-to-end UE channel. This is a limitation of 4G that 5G has ambitions to correct as part of its support for network slicing.

Support for mobility can now be understood as the process of re-executing one or more of the steps shown in [Figure 13](#) as the UE moves throughout the RAN. The unauthenticated link indicated by (1) allows the UE to be known to all base stations within range. (We refer to these as *potential links* in later chapters.) Based on the signal's measured CQI, the base stations communicate directly with each other to make a handover decision. Once made, the decision is then communicated to the Mobile Core, re-triggering the setup functions indicated by (3), which in turn re-builds the user plane tunnel between the base station and the Core-UP shown in [Figure 14](#). One of the most unique features of the cellular network is that the Mobile Core's user plane buffers data while idle UEs are transiting to active state, thereby avoiding dropped packets and subsequent end-to-end retransmissions.

In other words, the mobile cellular network maintains the *UE session* in the face of mobility (corresponding to the control and data channels depicted by (2) and (4) in [Figure 13](#), respectively), but it is able to do so only when the same Mobile Core serves the UE (i.e., only the base station changes). This would typically be the case for a UE moving within a metropolitan area. Moving between metro areas—and hence, between Mobile Cores—is indistinguishable from power cycling a UE. The UE is assigned a new IP address and no attempt is made to buffer and subsequently deliver in-flight data. Independent of mobility, but relevant to this discussion, any UE that becomes inactive for a period of time also loses its session, with a new session established and a new IP address assigned when the UE becomes active again.

Note that this session-based approach can be traced to the mobile cellular network's roots as a connection-oriented network. An interesting thought experiment is whether the Mobile Core will continue to evolve so as to better match the connectionless

assumptions of the Internet protocols that typically run on top of it.

We conclude this overview of the Mobile Core by returning to the role it plays in implementing a *global* mobile network. It is probably already clear that each MNO implements a database of subscriber information, allowing it to map an IMSI to a profile that records what services (roaming, data plane, hot spot support) the subscriber is paying for. This record also includes the international phone number for the device. How this database is realized is an implementation choice (of which we'll see an example in Chapter 5), but 3GPP defines an interface by which one Mobile Core (running on behalf of one MNO) queries another Mobile Core (running on behalf of some other MNO), to map between the IMSI, the phone number, and the subscriber profile.

2.5 Managed Cloud Service

The architectural overview presented up to this point focuses on the functional elements of the mobile cellular network. We now turn our attention to how this functionality is operationalized, and we do so in a decidedly software-defined and cloud-native way. This lays the foundation for the rest of the book, which builds towards supporting 5G connectivity as a managed cloud service. This is a marked change from the conventional Telco approach, whereby an operator bought purpose-built devices from a handful of vendors, and then managed them using the legacy OSS/BSS machinery that was originally designed for the telephony network.[\[1\]](#)

OSS/BSS stands for Operation Support System / Business Support System, and even traditional MNOs are now re-imagining them by adopting cloud practices. But this transition is a slow process due to all the legacy systems the Telcos need to continue supporting.

When we talk about “operationalizing” a network, we are referring to a substantial system that operators (whether they are traditional MNOs or cloud service providers) use to activate and manage all the constituent components (whether they are purpose-built devices or software running on commodity hardware). And because these network operators are people, one high-level summary is that this

management layer (whether it is an OSS/BSS or a cloud orchestrator) provides a way to map high-level *Intents* onto low-level *Actions*.

Figure 15. High-level summary of the role operationalization plays in a network deployment.

This overview, as illustrated in [Figure 15](#), is obviously quite abstract. To make the discussion more concrete, we use an open source implementation, called Aether, as an example. Aether is a Kubernetes-based edge cloud, augmented with a 5G-based connectivity service. Aether is targeted at enterprises that want to take advantage of 5G connectivity in support of edge applications that require predictable, low-latency connectivity. In short, “Kubernetes-based” means Aether is able to host container-based services, with Kubernetes being the platform used to orchestrate the services, and “5G-based connectivity” means Aether is able to connect those services to mobile devices throughout the enterprise’s physical environment.

Aether supports this combination by implementing both the RAN and the user plane of the Mobile Core on-prem, as cloud-native workloads co-located on the Aether cluster. This is often referred to as *local breakout* because it enables direct communication between mobile devices and edge applications without data traffic leaving the enterprise, in contrast to what would happen with standard, operator-provided 5G service. This scenario is depicted in [Figure 16](#), which shows unnamed edge applications running on-prem. Those edge applications might include the local processing of sensor data or control applications for the IoT devices, for example.


Figure 16. Overview of Aether as a hybrid cloud, with edge apps and the 5G data plane (called *local breakout*) running on-prem and various management and control-related workloads running in a central cloud.

The approach includes both edge (on-prem) and centralized (off-prem) components. This is true for edge apps, which often have a centralized counterpart running in a commodity cloud. It is also true for the 5G Mobile Core, where the on-prem User Plane (UP) is paired with a centralized Control Plane (CP). The central cloud


shown in this figure might be private (i.e., operated by the enterprise), public (i.e., operated by a commercial cloud provider), or some combination of the two (i.e., not all centralized elements need to run in the same cloud).

Also shown in [Figure 16](#) is a centralized *Control and Management Platform*. This is Aether's version of the "Management Layer" depicted in [Figure 15](#), and it represents all the functionality needed to offer Aether as a managed cloud service, with system administrators using a portal exported by this platform to operate the underlying infrastructure and services within their enterprise.


Once we deconstruct the individual components in more details in the next three chapters, we return to the question of how the resulting set of components can be assembled into an operational edge cloud in Chapter 6. The end result is 5G connectivity as a managed cloud service.



[1]



[2]



Chapter 3: Radio

Transmission

For anyone familiar with wireless access technologies that offer a best-effort service, such as Wi-Fi, the cellular network presents a notable contrast. This is mostly because cellular networks carefully allocate available radio spectrum among many users (or more precisely, UEs), aiming to deliver a certain quality to all active users in a coverage area, while also allowing those users to remain connected while moving. They also aim to maximize the efficiency of spectrum usage, as it is a finite and often costly resource. This has resulted in a highly dynamic and adaptive approach, in which coding, modulation and scheduling play a central role. 5G takes the cellular approach to managing radio resources to a new level of sophistication.

Wireless transmission is full of challenges that don't arise in wired networks. Interference can arise from many sources ranging from microwave ovens to baby monitors, while radio signals reflect off objects such as buildings and vehicles. Some objects absorb wireless signals. The properties of the wireless channel vary over time, and depending on the frequency in use. Much of the design of cellular radio systems is motivated by the need to deal with these challenges.

As we will see in this chapter, mobile cellular networks use a reservation-based strategy, whereas Wi-Fi is contention-based. This difference is rooted in each system's fundamental assumption about utilization: Wi-Fi assumes a lightly loaded network (and hence optimistically transmits when the wireless link is idle and backs off if contention is detected), while 5G assumes (and strives for) high utilization (and hence explicitly assign different users to different "shares" of the available radio spectrum).

The fact that 5G controls spectrum allocation carefully is central to its ability to deliver guarantees of bandwidth and latency more predictably than Wi-Fi. This in turn is why there is so much interest in using 5G for mission-critical applications.

We start by giving a short primer on radio transmission as a way of laying a foundation for understanding the rest of the 5G architecture. The following is not a substitute for a theoretical treatment of the topic, but is instead intended as a way of grounding the systems-oriented description of 5G that follows in the reality of wireless communication.

3.1 Coding and Modulation

The mobile channel over which digital data needs to be reliably transmitted brings a number of impairments, including noise, attenuation, distortion, fading, and interference. This challenge is addressed by a combination of coding and modulation, as depicted in [Figure 17](#).

Figure 17. The role of coding and modulation in mobile communication.

At its core, coding inserts extra bits into the data to help recover from all the environmental factors that interfere with signal propagation. This typically implies some form of *Forward Error Correction* (e.g., turbo codes, polar codes). Modulation then generates signals that represent the encoded data stream, and it does so in a way that matches the channel characteristics: It first uses a digital modulation signal format that maximizes the number of reliably transmitted bits every second based on the specifics of the observed channel impairments; it next matches the transmission bandwidth to channel bandwidth using pulse shaping; and finally, it uses RF modulation to transmit the signal as an electromagnetic wave over an assigned *carrier frequency*.

For a deeper appreciation of the challenges of reliably transmitting data by propagating radio signals through the air, consider the scenario depicted in [Figure 18](#), where the signal bounces off various stationary and moving objects, following multiple paths from the transmitter to the receiver, who may also be moving.

Figure 18. Signals propagate along multiple paths from transmitter to receiver.

As a consequence of these multiple paths, the original signal arrives at the receiver spread over time, as illustrated in [Figure 19](#). Empirical evidence shows that the Multipath Spread—the time between the first and last signals of one transmission arriving at the receiver—is 1-10 μ s in urban environments and 10-30 μ s in suburban environments. These multipath signals can interfere with each other constructively or destructively, and this will vary over time. Theoretical bounds for the time duration for which the channel may be assumed to be invariant, known as the *Coherence Time* and denoted T_c , is given by

$$T_c = c/v \times 1/f$$

where c is the velocity of the signal, v is the velocity of the receiver (e.g., moving car or train), and f is the frequency of the carrier signal that is being modulated. This says the coherence time is inversely proportional to the frequency of the signal and the speed of movement, which makes intuitive sense: The higher the frequency (narrower the wave) the shorter the coherence time, and likewise, the faster the receiver is moving the shorter the coherence time. Based on the target parameters to this model (selected according to the target physical environment), it is possible to calculate T_c , which in turn bounds the rate at which symbols can be transmitted without undue risk of interference. The dynamic nature of the wireless channel is a central challenge to address in the cellular network.

Figure 19. Received data spread over time due to multipath variation.

To complicate matters further, [Figure 18](#) and [19](#) imply the transmission originates from a single antenna, but cell towers are equipped with an array of antennas, each transmitting in a different (but overlapping) direction. This technology, called *Multiple-Input-Multiple-Output (MIMO)*, opens the door to purposely transmitting data from multiple antennas in an effort to reach the receiver, adding even more paths to the environment-imposed multipath propagation.

One of the most important consequences of these factors is that the transmitter must receive feedback from every receiver to judge how

to best utilize the wireless medium on their behalf. 3GPP specifies a *Channel Quality Indicator (CQI)* for this purpose. In practice, the receiver sends a CQI status report to the base station periodically (e.g., every millisecond). These CQI messages report the observed signal-to-noise ratio, which impacts the receiver's ability to recover the data bits. The base station then uses this information to adapt how it allocates the available radio spectrum to the subscribers it is serving, as well as which coding and modulation scheme to employ. All of these decisions are made by the scheduler.

3.2 Scheduler

How the scheduler does its job is one of the most important properties of each generation of the cellular network, which in turn depends on the multiplexing mechanism. For example, 2G used *Time Division Multiple Access (TDMA)* and 3G used *Code Division Multiple Access (CDMA)*. How data is multiplexed is also a major differentiator for 4G and 5G, completing the transition from the cellular network being fundamentally circuit-switched to fundamentally packet-switched.

Both 4G and 5G are based on *Orthogonal Frequency-Division Multiplexing (OFDM)*, an approach that multiplexes data over multiple orthogonal subcarrier frequencies, each of which is modulated independently. One attraction of OFDM is that, by splitting the frequency band into subcarriers, it can send symbols on each subcarrier at a relatively low rate. This makes it easier to correctly decode symbols in the presence of multipath interference. The efficiency of OFDM depends on the selection of subcarrier frequencies so as to avoid interference, that is, how it achieves orthogonality. That topic is beyond the scope of this book.

As long as you understand that OFDM uses a set of subcarriers, with symbols (each of which contain a few bits of data) being sent at some rate on each subcarrier, that you can appreciate that there are discrete schedulable units of the radio spectrum. The fundamental unit is the time to transmit one symbol on one subcarrier. With that building block in mind, we are now in a position to examine how multiplexing and scheduling work in 4G and 5G.

3.2.1 Multiplexing in 4G

We start with 4G because it provides a foundational understanding that makes 5G easier to explain, where both 4G and 5G use an approach to multiplexing called *Orthogonal Frequency-Division Multiple Access (OFDMA)*. You can think of OFDMA as a specific application of OFDM. In the 4G case, OFDMA multiplexes data over a set of 12 orthogonal (non-interfering) subcarrier frequencies, each of which is modulated independently. [1] The “Multiple Access” in OFDMA implies that data can simultaneously be sent on behalf of multiple users, each on a different subcarrier frequency and for a different duration of time. The 4G-defined subbands are narrow (e.g., 15 kHz), and the coding of user data into OFDMA symbols is designed to minimize the risk of data loss due to interference.

4G uses OFDMA for downstream traffic, and a different multiplexing strategy for upstream transmissions (from user devices to base stations), but we do not describe it because the approach is not applicable to 5G.

The use of OFDMA naturally leads to conceptualizing the radio spectrum as a 2-D resource, as shown in [Figure 20](#), with the subcarriers represented in the vertical dimension and the time to transmit symbols on each subcarrier represented in the horizontal dimension. The basic unit of transmission, called a *Resource Element (RE)*, corresponds to a 15-kHz band around one subcarrier frequency and the time it takes to transmit one OFDMA symbol. The number of bits that can be encoded in each symbol depends on the modulation scheme in use. For example, using *Quadrature Amplitude Modulation (QAM)*, 16-QAM yields 4 bits per symbol and 64-QAM yields 6 bits per symbol. The details of the modulation need not concern us here; the key point is that there is a degree of freedom to choose the modulation scheme based on the measured channel quality, sending more bits per symbol (and thus more bits per second) when the quality is high.

Figure 20. Spectrum abstractly represented by a 2-D grid of schedulable Resource Elements.

A scheduler allocates some number of REs to each user that has

data to transmit during each 1 ms *Transmission Time Interval (TTI)*, where users are depicted by different colored blocks in [Figure 20](#). The only constraint on the scheduler is that it must make its allocation decisions on blocks of $7 \times 12 = 84$ resource elements, called a *Physical Resource Block (PRB)*. [Figure 20](#) shows two back-to-back PRBs. Of course time continues to flow along one axis, and depending on the size of the available frequency band (e.g., it might be 100 MHz wide), there may be many more subcarrier slots (and hence PRBs) available along the other axis, so the scheduler is essentially preparing and transmitting a sequence of PRBs.

Note that OFDMA is not a coding/modulation algorithm, but instead provides a framework for selecting a specific coding and modulation for each subcarrier frequency. QAM is one common example modulation. It is the scheduler's responsibility to select the modulation to use for each PRB, based on the CQI feedback it has received. The scheduler also selects the coding on a per-PRB basis, for example, by how it sets the parameters to the turbo code algorithm.

The 1-ms TTI corresponds to the time frame in which the scheduler receives feedback from users about the quality of the signal they are experiencing. This is the role of CQI: once every millisecond, each user sends a set of metrics, which the scheduler uses to make its decision as to how to allocate PRBs during the subsequent TTI.

Another input to the scheduling decision is the *QoS Class Identifier (QCI)*, which indicates the quality-of-service each class of traffic is to receive. In 4G, the QCI value assigned to each class (there are twenty six such classes, in total) indicates whether the traffic has a *Guaranteed Bit Rate (GBR)* or not (*non-GBR*), plus the class's relative priority within those two categories. (Note that the 5QI parameter introduced in Chapter 2 serves the same purpose as the QCI parameter in 4G.)

Finally, keep in mind that [Figure 20](#) focuses on scheduling transmissions from a single antenna, but the MIMO technology described above means the scheduler also has to determine which antenna (or more generally, what subset of antennas) will most effectively reach each receiver. But again, in the abstract, the scheduler is charged with allocating a sequence of Resource

Elements.

Note that the scheduler has many degrees of freedom: it has to decide which set of users to service during a given time interval, how many resource elements to allocate to each such user, how to select the coding and modulation levels, and which antenna to transmit their data on. This is an optimization problem that, fortunately, we are not trying to solve here. Our goal is to describe an architecture that allows someone else to design and plug in an effective scheduler. Keeping the cellular architecture open to innovations like this is one of our goals, and as we will see in the next section, becomes even more important in 5G where the scheduler operates with even more degrees of freedom.

3.2.2 Multiplexing in 5G

The transition from 4G to 5G introduces additional flexibility in how the radio spectrum is scheduled, making it possible to adapt the cellular network to a more diverse set of devices and application domains.

Fundamentally, 5G defines a family of waveforms—unlike LTE, which specified only one waveform—each optimized for a different band in the radio spectrum.[\[2\]](#) The bands with carrier frequencies below 1 GHz are designed to deliver mobile broadband and massive IoT services with a primary focus on range. Carrier frequencies between 1-6 GHz are designed to offer wider bandwidths, focusing on mobile broadband and mission-critical applications. Carrier frequencies above 24 GHz (mmWaves) are designed to provide super-wide bandwidths over short, line-of-sight coverage.

A waveform is defined by the frequency, amplitude, and phase-shift independent property (shape) of a signal. A sine wave is a simple example of a waveform.

These different waveforms affect the scheduling and subcarrier intervals (i.e., the “size” of the resource elements described in the previous section).

- For frequency range 1 (410 MHz - 7125 MHz), 5G allows maximum 100 MHz bandwidths. In this case, there are three waveforms with subcarrier spacings of 15, 30 and 60 kHz. (We used 15 kHz in the example shown in [Figure 20](#).) The corresponding scheduling intervals of 0.5, 0.25, and 0.125 ms, respectively. (We used 0.5 ms in the example shown in [Figure 20](#).)
- For millimeter bands, also known as frequency range 2 (24.25 GHz - 52.6 GHz), bandwidths may go from 50 MHz up to 400 MHz. There are two waveforms, with subcarrier spacings of 60 kHz and 120 kHz. Both have scheduling intervals of 0.125 ms.

These various configurations of subcarrier spacing and scheduling intervals are sometimes called the *numerology* of the radio's air interface.

This range of numerology is important because it adds another degree of freedom to the scheduler. In addition to allocating radio resources to users, it has the ability to dynamically adjust the size of the resource by changing the waveform being used. With this additional freedom, fixed-sized REs are no longer the primary unit of resource allocation. We instead use more abstract terminology, and talk about allocating *Resource Blocks* to subscribers, where the 5G scheduler determines both the size and number of Resource Blocks allocated during each time interval.

[Figure 21](#) depicts the role of the scheduler from this more abstract perspective. Just as with 4G, CQI feedback from the receivers and the 5QI quality-of-service class selected by the subscriber are the two key pieces of input to the scheduler. Note that the set of 5QI values available in 5G is considerably greater than its QCI counterpart in 4G, reflecting the increasing differentiation among classes that 5G aims to support. For 5G, each class includes the following attributes:

- Resource Type: Guaranteed Bit Rate (GBR), Delay-Critical GBR, Non-GBR
- Priority Level
- Packet Delay Budget
- Packet Error Rate

- Maximum Data Burst
- Averaging Window

Note that while the preceding discussion could be interpreted to imply a one-to-one relationship between subscribers and a 5QI, it is more accurate to say that each 5QI is associated with a class of traffic (often corresponding to some type of application), where a given subscriber might be sending and receiving traffic that belongs to multiple classes at any given time.

Figure 21. Scheduler allocates Resource Blocks to user data streams based on CQI feedback from receivers and the 5QI parameters associated with each class of service.

3.3 Virtualized Scheduler (Slicing)

The discussion up to this point presumes a single scheduler is suitable for all workloads, but different applications have different requirements for how their traffic gets scheduled. For example, some applications care about latency and others care more about bandwidth.

While in principle one could define a sophisticated scheduler that takes dozens of different factors into account, 5G has introduced a mechanism that allows the underlying resources (in this case radio spectrum) to be “sliced” between different uses. The key to slicing is to add a layer of indirection between the scheduler and the physical resource blocks. Slicing, like much of 5G, has received a degree of hype, but it boils down to virtualization at the level of the radio scheduler.

As shown in [Figure 22](#), the idea is to first schedule offered traffic demands to virtual RBs, and then perform a second mapping of Virtual RBs to Physical RBs. This sort of virtualization is common in resource allocators throughout computing systems because we want to separate how many resources are allocated to each user (or virtual machine in the computing case) from the decision as to

which physical resources are actually assigned. This virtual-to-physical mapping is performed by a layer typically known as a *Hypervisor*, and the important thing to keep in mind is that it is totally agnostic about which user's segment is affected by each translation.

Figure 22. Wireless Hypervisor mapping virtual resource blocks to physical resource blocks.

Having decoupled the Virtual RBs from Physical RBs, it is now possible to define multiple Virtual RB sets (of varying sizes), each with its own scheduler. [Figure 23](#) gives an example with two equal-sized RB sets. The important consequence is this: having made the macro-decision that the Physical RBs are divided into two equal partitions, the scheduler associated with each partition is free to allocate Virtual RBs independently from the other. For example, one scheduler might be designed to deal with high-bandwidth video traffic and another scheduler might be optimized for low-latency IoT traffic. Alternatively, a certain fraction of the available capacity could be reserved for premium customers or other high-priority traffic (e.g., public safety), with the rest shared among everyone else.

Figure 23. Multiple schedulers running on top of wireless hypervisor.

A final point to note is that there is considerable flexibility in the allocation of resources to slices. While the example above shows resources allocated in a fixed manner to each slice, it is possible to make unused resources in one slice available to another slice, as long as they can be reclaimed when needed. This is similar to how work-conserving scheduling works in network queues: resources are not wasted if the offered load in a class is less than the rate guaranteed for that class.

3.4 New Use Cases

We conclude by noting that up to this point we have described 5G as introducing additional degrees of freedom into how data is

scheduled for transmission, but when taken as a whole, the end result is a qualitatively more powerful radio. This new 5G air interface specification, which is commonly referred to as *New Radio (NR)*, enables new use cases that go well beyond simply delivering increased bandwidth. 3GPP defined five such use cases:

- Enhanced Mobile Broadband (eMBB)
- Ultra-Reliable Low-Latency Communications (URLLC)
- Massive Internet of Things (MIoT)
- Vehicle to Infrastructure/Vehicle (V2X)
- High-Performance Machine-Type Communications (HMTC)

These use cases reflect the requirements introduced in Chapter 1, and can be attributed to four fundamental improvements in how 5G multiplexes data onto the radio spectrum.

The first improvement is being able to change the waveform. This effectively introduces the ability to dynamically change the size and number of schedulable resource units, which opens the door to making fine-grained scheduling decisions that are critical to predictable, low-latency communication.

The second is related to the “Multiple Access” aspect of how distinct traffic sources are multiplexed onto the available spectrum. In 4G, multiplexing happens in both the frequency and time domains for downstream traffic, but only in the frequency domain for upstream traffic. 5G NR multiplexes both upstream and downstream traffic in both the time and frequency domains. Doing so provides finer-grained scheduling control needed by latency-sensitive applications.


The third advance is related to the new spectrum available to 5G NR, with the mmWave allocations opening above 24 GHz being especially important. This is not only because of the abundance of capacity—which makes it possible to set aside dedicated capacity for applications that require low-latency communication—but also because the higher frequency enables even finer-grained resource blocks (e.g., scheduling intervals as short as 0.125 ms). Again, this improves scheduling granularity to the benefit of applications that cannot tolerate unpredictable latency.

The fourth is related to delivering mobile connectivity to a massive



number of IoT devices, ranging from devices that require mobility support and modest data rates (e.g. wearables, asset trackers) to devices that support intermittent transmission of a few bytes of data (e.g., sensors, meters). None of these devices are particularly latency-sensitive or bandwidth-hungry, but they often require long battery lifetimes, and hence, reduced hardware complexity that draws less power.



Support for IoT device connectivity revolves around allocating some of the available radio spectrum to a light-weight (simplified) air interface. This approach started with Release 13 of LTE via two complementary technologies: mMTC and NB-IoT (NarrowBand-IoT). Both technologies build on a significantly simplified version of LTE—i.e., limiting the numerology and flexibility needed to achieve high spectrum utilization—so as to allow for simpler IoT hardware design. mMTC delivers up to 1 Mbps over 1.4 MHz of bandwidth and NB-IoT delivers a few tens of kbps over 200 kHz of bandwidth; hence the term *NarrowBand*. Both technologies have been designed to support over 1 million devices per square kilometer. With Release 16, both technologies can be operated in-band with 5G, but still based on LTE numerology. Starting with Release 17, a simpler version of 5G NR, called *NR-Light*, will be introduced as the evolution of mMTC. NR-Light is expected to scale the device density even further.

As a complement to these four improvements, 5G NR is designed to support partitioning the available bandwidth, with different partitions dynamically allocated to different classes of traffic (e.g., high-bandwidth, low-latency, and low-complexity). This is the essence of *slicing*, as discussed above. Moreover, once traffic with different requirements can be served by different slices, 5G NR's approach to multiplexing is general enough to support varied scheduling decisions for those slices, each tailored for the target traffic. We return to the applications of slicing in Chapter 6.



Other Base Stations
(for Handover or
Link Aggregation)





Example
xApps

Handover Control

Interference Mgmt

Link Aggregation

Load Balancing

..


RIC

R-NIB

CU


DU

DU



RU


RU




RU


RU

[1]



[2]





Chapter 4: Radio Access Network

The high-level description of the RAN in Chapter 2 was mostly silent about the RAN's internal structure. We now focus on those details, and in doing so, explain how the RAN is being transformed in 5G.

You can think of the RAN as having one main job: to transfer packets between the mobile core and a set of UEs. This means it is deeply involved in the management and scheduling of radio spectrum that we discussed in the last chapter, but there is more to it than that. We start by describing the stages in the RAN's packet processing pipeline, and then showing how these stages are being disaggregated, distributed, and implemented.

Note that the deconstruction of the RAN presented in this chapter represents a combination of standardized specifications and implementation strategies. The former continues to be under the purview of the 3GPP, but the latter are primarily influenced by a second organization: the *Open-RAN Alliance (O-RAN)* introduced in Chapter 1. O-RAN is led by network operators with the goal of developing a software-based implementation of the RAN that

eliminates vendor lock-in. Such business forces are certainly a factor in where 5G mobile networks are headed, but our goal in this chapter is to identify the technical design decisions involved in that evolution.

4.1 Packet Processing Pipeline

Figure 24 shows the packet processing stages historically bundled in base stations, as specified by the 3GPP standard. Note that the figure depicts the base station as a pipeline (running left-to-right for packets sent to the UE) but it is equally valid to view it as a protocol stack (as is typically done in official 3GPP documents). Also note that (for now) we are agnostic as to how these stages are implemented. Since we are ultimately heading towards a cloud-based implementation, one possible implementation strategy would be a microservice per box.

Figure 24. RAN processing pipeline, including both user and control plane components.

The key stages are as follows.

- RRC (Radio Resource Control) → Responsible for configuring the coarse-grained and policy-related aspects of the pipeline. The RRC runs in the RAN's control plane; it does not process packets in the user plane.
- PDCP (Packet Data Convergence Protocol) → Responsible for compressing and decompressing IP headers, ciphering and deciphering, integrity protection and integrity verification, duplication, reordering and in-order delivery, and out-of-order delivery.
- RLC (Radio Link Control) → Responsible for segmentation and reassembly, as well as reliably transmitting/receiving segments using error correction through ARQ (automatic repeat request).
- MAC (Media Access Control) → Responsible for buffering, multiplexing and demultiplexing segments, including all real-

time scheduling decisions about what segments are transmitted when. Also able to make a “late” forwarding decision (i.e., to alternative carrier frequencies, including Wi-Fi).

- PHY (Physical Layer) → Responsible for coding and modulation (as discussed in Chapter 3), including FEC.

The last two stages in [Figure 24](#) (D/A conversion and the RF front-end) are beyond the scope of this book.

While it is simplest to view the stages in [Figure 24](#) as a pure left-to-right pipeline, the Scheduler described in Section 3.2 (denoted “S” in the figure) runs in the MAC stage/layer, and implements the “main loop” for outbound traffic: It reads data from the upstream RLC and schedules transmissions to the downstream PHY. Since the Scheduler determines the number of bytes to transmit to a given UE during each time period (based on all the factors outlined in Chapter 3), it must request (get) a segment of that length from the upstream queue. In practice, the size of the segment that can be transmitted on behalf of a single UE during a single scheduling interval can range from a few bytes to an entire IP packet.

Also note that a combination of the RRC and PDCP are responsible for the observation made in Section 2.3: that a “*base station can be viewed as a specialized forwarder*”. The control plane logic that decides whether this base station should continue processing a packet or forward it to another base station runs in the RRC, and the corresponding user plane mechanism that carries out the forwarding decision is implemented in the PDCP. The interface between these two elements is defined by the 3GPP spec, but the decision-making logic is an implementation choice (and historically proprietary). This control logic is generically referred as the *Radio Resource Management (RRM)*, not to be confused with the standards-defined RRC stage depicted in [Figure 24](#).

4.2 Split RAN

The next step is to understand how the functionality outlined above is partitioned between physical elements, and hence, “split” across

centralized and distributed locations. The dominant option has historically been “no split,” with the entire pipeline shown in [Figure 24](#) running in the base station. Going forward, the 3GPP standard has been extended to allow for multiple split-points, with the partition shown in [Figure 25](#) being actively pursued by the operator-led O-RAN Alliance. It is the split we adopt throughout the rest of this book. Note that the split between centralized and distributed components mirrors the split made in SDN, with similar motivations. We discuss further how SDN techniques are applied to the RAN below.

Figure 25. Split RAN processing pipeline distributed across a Central Unit (CU), Distributed Unit (DU), and Radio Unit (RU).

This results in a RAN-wide configuration similar to that shown in [Figure 26](#), where a single *Central Unit (CU)* running in the cloud serves multiple *Distributed Units (DUs)*, each of which in turn serves multiple *Radio Units (RUs)*. Critically, the RRC (centralized in the CU) is responsible for making only near-real-time configuration and control decisions, while the Scheduler that is part of the MAC stage is responsible for all real-time scheduling decisions.

Figure 26. Split RAN hierarchy, with one CU serving multiple DUs, each of which serves multiple RUs.

Because scheduling decisions for radio transmission are made by the MAC layer in real time, a DU needs to be “near” (within 1ms) the RUs it manages. (You can’t afford to make scheduling decisions based on out-of-date channel information.) One familiar configuration is to co-locate a DU and an RU in a cell tower. But when an RU corresponds to a small cell, many of which might be spread across a modestly-sized geographic area (e.g., a mall, campus, or factory), then a single DU would likely service multiple RUs. The use of mmWave in 5G is likely to make this latter configuration all the more common.

Also note that the Split RAN changes the nature of the Backhaul Network, which originally connected the base stations back to the Mobile Core. With the Split RAN there are multiple connections, which are officially labeled as follows.

- RU-DU connectivity is called the Fronthaul
- DU-CU connectivity is called the Midhaul
- CU-Mobile Core connectivity is called the Backhaul

For more insight into design considerations for interconnecting the distributed components of a Split RAN, we recommend the NGMN Alliance Report.

Further Reading

[RAN Evolution Project: Backhaul and Fronthaul Evolution](#). NGMN Alliance Report, March 2015.

One observation about the CU (which becomes relevant in Chapter 6 when we incorporate it into a managed cloud service) is that one might co-locate the CU and Mobile Core in the same cluster, meaning the backhaul is implemented in the cluster switching fabric. In such a configuration, the midhaul then effectively serves the same purpose as the original backhaul, and the fronthaul is constrained by the predictable/low-latency requirements of the MAC stage's real-time scheduler. This situation is further complicated by the fact that the mobile core itself may be disaggregated, as discussed in Chapter 5.

A second observation about the CU shown in [Figure 25](#) is that it encompasses two functional blocks—the RRC and the PDCP—which lie on the RAN's control plane and user plane, respectively. This separation is consistent with the idea of CUPS introduced in Chapter 2, and plays an increasingly important role as we dig deeper into how the RAN is implemented. For now, we note that the two parts are sometimes referred to as the CU-CP and CU-UP, respectively.

We conclude our description of the split RAN architecture with the alternative depiction in [Figure 27](#). For completeness, this figure identifies the standardized interfaces between the components (e.g., N2, N3, F1-U, F1-C, and Open Fronthaul). We're not going to talk about these interfaces, except to note that they exist and there is a corresponding 3GPP specification that spells out the details. Instead, we're going to comment on the availability of open source implementations for each component.

Figure 27. Alternative depiction of the Split RAN components, showing the 3GPP-specified inter-unit interfaces.

With respect to the Central Unit, most of the complexity is in the CU-CP, which, as we'll see in the next section, is being re-engineered using SDN, with open source solutions available. With respect to the Radio Unit, nearly all the complexity is in D/A conversion and how the resulting analog signal is amplified. Incumbent vendors have significant proprietary know-how in this space, which will almost certainly remain closed.

With respect to the Distributed Unit, the news is mixed, and correspondingly, the figure shows more detail. The High-PHY module—which corresponds to all but the RF modulation step of [Figure 17](#) in Section 3.1—is one of the most complex components in the RAN stack. An open source implementation of the High-PHY, known as FlexRAN, exists and is widely used in commercial products. The only caveat is that the software license restricts usage to Intel processors, although it is also the case that the FlexRAN software exploits Intel-specific hardware capabilities. As for the rest of the DU, the MAC is the other source of high-value closed technology, particularly in how scheduling is done. There is an open source version made available by the Open Air Initiative (OAI), but its usage is restricted to research-only deployments.

Further Reading

[FlexRAN: Reference Architecture for Wireless Access.](#)

[Open Air Interface.](#)

4.3 Software-Defined RAN

We now describe how the RAN is being implemented according to SDN principles, resulting in an SD-RAN. The key architectural insight is shown in [Figure 28](#), where the RRC from [Figure 24](#) is partitioned into two sub-components: the one on the left provides a 3GPP-compliant way for the RAN to interface to the Mobile Core's control plane (the figure labels this sub-component as a "Proxy"), while the one on the right opens a new programmatic API for

exerting software-based control over the pipeline that implements the RAN user plane.

To be more specific, the left sub-component simply forwards control packets between the Mobile Core and the PDCP, providing a path over which the Mobile Core can communicate with the UE for control purposes, whereas the right sub-component implements the core of the RRC's control functionality (which as we explained in Section 4.1 is also known as RRM). This latter component is commonly referred to as the *RAN Intelligent Controller (RIC)* in O-RAN architecture documents, so we adopt this terminology. The “Near-Real Time” qualifier indicates the RIC is part of 10-100ms control loop implemented in the CU, as opposed to the ~1ms control loop required by the MAC scheduler running in the DU.

Figure 28. RRC disaggregated into a Mobile Core facing control plane component (a proxy) and a Near-Real-Time Controller.

Although not shown in [Figure 28](#), keep in mind (from [Figure 25](#)) that the RRC and the PDCP form the CU. Reconciling these two figures is a little bit messy, but to a first approximation, the PDCP corresponds to the CU-UP and RRC-Proxy corresponds to the CU-CP, with the RIC “lifted out” and responsible for overseeing both. We postpone a diagram depicting this relationship until Section 4.5, where we summarize the end-to-end result. For now, the important takeaway is that the SDN-inspired refactoring of the RAN is free both to move functionality around and to introduce new module boundaries, as long as the original 3GPP-defined interfaces are preserved.

Figure 29. Example set of control applications (xApps) running on top of Near-Real-Time RAN Controller (RIC), controlling a distributed set of Split RAN elements (CU, DU, RU).

Completing the picture, [Figure 29](#) shows the Near-RT RIC implemented as an SDN Controller hosting a set of SDN control apps. The RIC maintains a *RAN Network Information Base (R-NIB)*—a common set of information that can be consumed by numerous control apps. The R-NIB includes time-averaged CQI values and other per-session state (e.g., GTP tunnel IDs, 5QI values for the type of traffic), while the MAC (as part of the DU) maintains the

instantaneous CQI values required by the real-time scheduler. More generally, the R-NIB includes the following state:

- Fixed Nodes (RU/DU/CU Attributes)
 - Identifiers
 - Version
 - Config Report
 - RRM config
 - PHY resource usage
- Mobile Nodes (UE Attributes)
 - Devices
 - Identifiers
 - Capability
 - Measurement Config
 - State (Active/Idle)
 - Links (*Actual* and *Potential*)
 - Identifiers
 - Link Type
 - Config/Bearer Parameters
 - 5QI Value
- Virtual Constructs (Slice Attributes)
 - Links
 - Bearers/Flows
 - Validity Period
 - Desired KPIs
 - MAC RRM Configuration
 - RRM Control Configuration

The four example Control Apps (xApps) in [Figure 29](#) do not constitute an exhaustive list, but they do represent the sweet spot for SDN, with its emphasis on central control over distributed forwarding. These functions—Link Aggregation Control, Interference Management, Load Balancing, and Handover Control—are often implemented by individual base stations with only local visibility, but they have global consequences. The SDN approach is to collect the available input data centrally, make a globally optimal decision, and then push the respective control parameters back to the base stations for execution. Evidence using an analogous approach to optimize wide-area networks over many years (see, for example, B4) is compelling.

Further Reading

For an example of how SDN principles have been successfully applied to a production network, we recommend [B4: Experience with a Globally-Deployed Software Defined WAN](#). ACM SIGCOMM, August 2013.

One way to characterize xApps is based on the current practice of controlling the mobile link at two different levels. At a fine-grained level, per-node and per-link control are conducted using the RRM functions that are distributed across the individual base stations. [1] RRM functions include scheduling, handover control, link and carrier aggregation control, bearer control, and access control. At a coarse-grained level, regional mobile network optimization and configuration is conducted using *Self-Organizing Network (SON)* functions. These functions oversee neighbor lists, manage load balancing, optimize coverage and capacity, aim for network-wide interference mitigation, manage centrally configured parameters, and so on. As a consequence of these two levels of control, it is not uncommon to see reference to *RRM Applications* and *SON Applications*, respectively, in O-RAN documents for SD-RAN. For example, the Interference Management and Load Balancing xApps in [Figure 29](#) are SON Applications, while the other two xApps are RRM Applications.

Pedantically, Radio Resource Management (RRM) is another name for the collection of control functionality typically implemented in the RRC stage of the RAN pipeline.

Keep in mind, however, that this characterization of xApps is based on past (pre-SDN) implementations of the RAN. This is helpful as the industry transitions to SD-RAN, but the situation is likely to change. SDN is transforming the RAN, so new ways of controlling the RAN—resulting in applications that do not fit neatly into the RRM vs SON dichotomy—can be expected to emerge over time.

4.4 Near Real-Time RIC

Drilling down to the next level of detail, [Figure 30](#) shows an

exemplar implementation of a RIC based on a retargeting of the Open Network OS (ONOS) for the SD-RAN use case. ONOS (described in our SDN book) was originally designed to support traditional wireline network switches using standard interfaces (OpenFlow, P4Runtime, etc.). For the SD-RAN use case, the ONOS-based RIC instead supports a set of RAN-specific north- and south-facing interfaces, but internally takes advantage of the same collection of subsystems (microservices) as in the wireline case.[\[2\]](#)

Technically, the O-RAN definition of the RIC refers to the combination of xApps and the underlying platform (in our case ONOS), but we emphasize the distinction between the two, in keeping with the SDN model of distinguishing between the Network OS and the suite of Control Apps that run on it.

Figure 30. O-RAN compliant RAN Intelligent Controller (RIC) built by adapting and extending ONOS.

Specifically, the ONOS-based RIC includes a Topology Service to keep track of the fixed RAN infrastructure, a Device Service to track the UEs, and a Configuration Service to manage RAN-wide configuration state. All three of these services are implemented as Kubernetes-based microservices, and take advantage of a scalable key-value store.

Of the three interfaces called out in [Figure 30](#), the **A1** and **E2** interfaces are based on pre-existing 3GPP standards. The third, **xApp SDK**, is specific to the ONOS-based implementation. The O-RAN Alliance is using it to drive towards a unified API (and corresponding SDK) for building RIC-agnostic xApps.

The A1 interface provides a means for the mobile operator's management plane—typically called the *OSS/BSS (Operations Support System / Business Support System)* in the Telco world—to configure the RAN. We briefly introduced the OSS/BSS in Section 2.5, but all you need to know about it for our purposes is that such a component sits at the top of all Telco software stacks. It is the source of all configuration settings and business logic needed to operate a network. You can think of A1 as the RAN's counterpart to gNMI/gNOI (gRPC Network Management Interface/gRPC Network Operations Interface), a pair of configuration APIs commonly used

to configure commodity cloud hardware.

The Near-RT RIC uses the E2 interface to control the underlying RAN elements, including the CU, DUs, and RUs. A requirement of the E2 interface is that it be able to connect the Near-RT RIC to different types of RAN elements from different vendors. This range is reflected in the API, which revolves around a *Service Model* abstraction. The idea is that each RAN element advertises a Service Model, which effectively defines the set of RAN Functions the element is able to support. The RIC then issues a combination of the following four operations against this Service Model.

- **Report:** RIC asks the element to report a function-specific value setting.
- **Insert:** RIC instructs the element to activate a user plane function.
- **Control:** RIC instructs the element to activate a control plane function.
- **Policy:** RIC sets a policy parameter on one of the activated functions.

Of course, it is the RAN element, through its published Service Model, that defines the relevant set of functions that can be activated, the variables that can be reported, and policies that can be set. The O-RAN community is working on three vendor-agnostic Service Models. The first, called *Key Performance Measurement* (abbreviated *E2SM-KPM*), specifies the metrics that can be retrieved from RAN elements. The second, called *RAN Control* (abbreviated *E2SM-RC*), specifies parameters that can be set to control RAN elements. The third, called *Cell Configuration and Control* (abbreviated *E2SM-CCC*), exposes configuration parameters at the cell level.

In simple terms, E2SM-KPM defines what values can be *read* and E2SM-RC and E2SM-CCC defines what values can be *written*. Because the available values can be highly variable across all possible devices, we can expect different vendors will support only a subset of the entire collection. This will limit the “universality” the O-RAN was hoping to achieve in an effort to break vendor lock-in, but that outcome is familiar to network operators who have been dealing with divergent *Management Information Bases (MIBs)*

since the earliest days of the Internet.

Finally, the xApp SDK, which is specific to the ONOS-based implementation, is currently little more than a “pass through” of the E2 interface. This implies the xApps are expected to be aware of the available Service Models. One of the challenges the SDK has to deal with is how data passed to/from the RAN elements is encoded. For historical reasons, the E2 interface uses ASN.1 formatting, whereas the ONOS-RIC internally uses gRPC and Protocol Buffers to communicate between the set of microservices. The south-bound E2 interface in [Figure 30](#) translates between the two formats. The SDK currently makes the gRPC-based API available to xApps.

Further Reading

To learn more about the details of ONOS and its interfaces, we recommend the chapter in our SDN book that covers it in depth. [Software-Defined Networks: A Systems Approach. Chapter 6: Network OS.](#)

4.5 Control Loops

We conclude this description of RAN internals by revisiting the sequence of steps involved in disaggregation, which, as the previous three sections reveal, is being pursued in multiple tiers. In doing so, we tie up several loose ends, and focus attention on the resulting three control loops.

In the first tier of disaggregation, 3GPP defines multiple options for how the RAN can be split and distributed, with the pipeline shown in [Figure 24](#) disaggregated into the independently operating CU, DU, and RU components shown in [Figure 31](#). The O-RAN Alliance has selected specific disaggregation options from 3GPP and is developing open interfaces between these components.

Figure 31. First tier of RAN disaggregation: Split RAN.

The second tier of disaggregation focuses on the control/user plane separation (CUPS) of the CU, resulting in the CU-UP and CU-CP shown in [Figure 32](#). The control plane in question is the 3GPP

control plane, where the CU-UP realizes a pipeline for user traffic and the CU-CP focuses on control message signaling between Mobile Core and the disaggregated RAN components (as well as to the UE).

Figure 32. Second tier of RAN disaggregation: CUPS.

The third tier follows the SDN paradigm by separating most of RAN control (RRC functions) from the disaggregated RAN components, and logically centralizing them as applications running on an SDN Controller, which corresponds to the Near-RT RIC shown previously in [Figures 28](#) and [29](#). This SDN-based disaggregation is repeated in [Figure 33](#), which also shows the O-RAN-prescribed interfaces A1 and E2 introduced in the previous section. (Note that all the edges in [Figures 31](#) and [32](#) correspond to 3GPP-defined interfaces, as identified in Section 4.2, but their details are outside the scope of this discussion.)



Figure 33. Third tier of RAN disaggregation: SDN.

Taken together, the A1 and E2 interfaces complete two of the three major control loops of the RAN: the outer (non-real-time) loop has the Non-RT RIC as its control point and the middle (near-real-time) loop has the Near-RT RIC as its control point. The third (innermost) control loop—shown in [Figure 33](#) running inside the DU—includes the real-time Scheduler embedded in the MAC stage of the RAN pipeline. The two outer control loops have rough time bounds of $> 1\text{sec}$ and $> 10\text{ms}$, respectively. As we saw in Chapter 2, the real-time control loop is assumed to be $< 1\text{ms}$.


This raises the question of how specific functionality is distributed between the Non-RT RIC, Near-RT RIC, and DU. Starting with the second pair (i.e., the two inner loops), it is the case that not all RRC functions can be centralized; some need to be implemented in the DU. The SDN-based disaggregation then focuses on those that can be centralized, with the Near-RT RIC supporting the RRC applications and the SON applications mentioned in Section 4.3.


Turning to the outer two control loops, the Near RT-RIC opens the possibility of introducing policy-based RAN control, whereby interrupts (exceptions) to operator-defined policies would signal the


need for the outer loop to become involved. For example, one can imagine developing learning-based controls, where the inference engines for these controls would run as applications on the Near RT-RIC, and their non-real-time learning counterparts would run elsewhere. The Non-RT RIC would then interact with the Near-RT RIC to deliver relevant operator policies from the Management Plane to the Near RT-RIC over the A1 interface.



[1]







Chapter 5: Mobile Core

The Mobile Core provides IP connectivity to the RAN. It authenticates UEs as they connect, tracks them as they move from one base station to another, ensures that this connectivity fulfills the promised QoS requirements, and meters usage for billing.

Historically, all of these functions were provided by one or more proprietary network appliances. But like the rest of the 5G mobile network, these appliances are being disaggregated and implemented as a set of cloud services, with the goal of improving feature

velocity for new classes of applications. It is also the case that as the range of use cases grows more diverse, a one-size-fits-all approach will become problematic. The expectation is that it should be possible to customize and specialize the Mobile Core on a per-application basis.

This chapter introduces the functional elements of the Mobile Core, and describes different strategies for implementing that functionality.

5.1 Identity Management

There are two equally valid views of the Mobile Core. The Internet-centric view is that each local instantiation of the Mobile Core (e.g., serving a metro area) acts as a router that connects a physical RAN (one of many possible access network technologies, not unlike WiFi) to the global Internet. In this view, IP addresses serve as the unique global identifier that makes it possible for any RAN-connected device to communicate with any Internet-addressable device or service. The 3GPP-centric view is that a distributed set of Mobile Cores (interconnected by one or more backbone technologies, of which the Internet is just one example) cooperate to turn a set of physical RANs into one logically global RAN. In this perspective, the IMSI burned into a device's SIM card serves as the global identifier that makes it possible for any two mobile devices to communicate with each other.

Both of these perspectives are correct, but since broadband communication using Internet protocols to access cloud services is the dominant use case, this section takes an Internet-centric perspective of the Mobile Core. But before getting to that, we first need to understand several things about the 3GPP-centric perspective.

For starters, we need to be aware of the distinction between “identity” and “identifier”. The first term is commonly used when talking about principals or users, and the second term is used when talking about abstract objects or physical devices. Unfortunately, the two terms are conflated in the 3GPP architecture: The acronym

IMSI explicitly includes the word “Identity”, where the “S” in both IMSI and SIM stands for subscriber (a kind of principal), yet the IMSI is also used as a global identifier for a UE connected to the mobile network. This conflation breaks down when there could be tens or hundreds of IoT devices for every customer, with no obvious association among them. Accounting for this problem is an “architecture alignment” fix we discuss in the next chapter when we describe how to provide Private 5G Connectivity as a managed cloud service.

If we take the view that an IMSI is primarily a global identifier for UEs, then we can think of it as the mobile network’s equivalent of a 48-bit 802.3 or 802.11 MAC address. This includes how addresses are assigned to ensure uniqueness: (MCC, MNC) pairs are assigned by a global authority to every MNO, each of which then decides how to uniquely assign the rest of the IMSI identifier space to devices. This approach is similar to how network vendors are assigned a unique prefix for all the MAC addresses they configure into the NIC cards and WiFi chips they ship, but with one big difference: It is the MNO, rather than the vendor, that is responsible for assigning IMSIs to SIM cards. This makes the IMSI allocation problem closer to how the Internet assigns IP addresses to end hosts, but unlike DHCP, the IMSI-to-device binding is static.

This is important because, unlike 802.11 addresses, IMSIs are also intended to support global routing. (Here, we are using a liberal notion of routing—to locate an object—and focusing on the original 3GPP-perspective of the global RAN in which the Internet is just a possible packet network that interconnects Mobile Cores.) A hierarchically distributed database maps IMSIs onto the collection of information needed to forward data to the corresponding device. This includes a combination of relatively *static* information about the level of service the device expects to receive (including the corresponding phone number and subscriber profile/account information), and more *dynamic* information about the current location of the device (including which Mobile Core, and which base station served by that Core, currently connects the device to the global RAN).

This mapping service has a name, or rather, several names that

keep changing from generation to generation. In 2G and 3G it was called HLR (Home Location Registry). In 4G the HLR maintains only static information and a separate HSS (Home Subscriber Server) maintains the more dynamic information. In 5G the HLR is renamed the UDR (Unified Data Registry) and the HSS is renamed UDM (Unified Data Management). We will see the UDM in Section 5.2 because of the role it plays *within* a single instance of the Mobile Core.

There are, of course, many more details to the process—including how to find a device that has roamed to another MNO’s network—but conceptually the process is straightforward. (As a thought experiment, imagine how you would build a “logically global WiFi” using just 802.11 addresses, rather than depending on the additional layer of addressing provided by IP.) The important takeaway is that IMSIs are used to locate the Mobile Core instance that is then responsible for authenticating the device, tracking the device as it moves from base station to base station within that Core’s geographic region, and forwarding packets to/from the device.

Two additional observations about the relationship between IMSIs and IP addresses are worth highlighting. First, the odds of someone trying to “call” or “text” an IoT device, drone, camera, or robot are virtually zero. It is the IP address assigned to each device (by the local Mobile Core) that is used to locate (route packets to) the device. In this context, the IMSI plays exactly the same role in a physical RAN as an 802.11 address plays in a LAN, and the Mobile Core behaves just like any access router.

Second, whether a device connects to a RAN or some other access network, it is automatically assigned a new IP address any time it moves from one coverage domain to another. Even for voice calls in the RAN case, ongoing calls are dropped whenever a device moves between instantiations of the Mobile Core (i.e., uninterrupted mobility is supported only within the region served by a given Core). This is typically not a problem when the RAN is being used to deliver broadband connectivity because Internet devices are almost always clients *requesting* a cloud service; they just start issuing requests with their new (dynamically assigned) IP address.

5.2 Functional Components

The 5G Mobile Core, which 3GPP calls the *5GC*, adopts a microservice-like architecture officially known as the 3GPP *Service Based Architecture*. We say “microservice-like” because while the 3GPP specification spells out this level of disaggregation, it is really just describing a set of functional blocks and not prescribing an implementation. In practice, a set of functional blocks is very different from the collection of engineering decisions that go into designing a microservice-based system. That said, viewing the collection of components shown in [Figure 34](#) as a set of microservices is a reasonable working model (for now).

Figure 34. 5G Mobile Core (5GC), represented as a collection of microservices, where 3GPP defines the interfaces connecting the Mobile Core CP and UP to the RAN (denoted N2 and N3, respectively).

Starting with the User Plane (UP), the *UPF (User Plane Function)* forwards traffic between the RAN and the Internet. In addition to IP packet forwarding, the UPF is responsible for policy enforcement, lawful intercept, traffic usage measurement, and QoS policing. These are all common functions in access routers, even if they go beyond what you usually find in enterprise or backbone routers. The other detail of note is that, because the RAN is an overlay network, the UPF is responsible for tunneling (i.e., encapsulating and decapsulating) packets as they are transmitted to and from base stations over the N3 interface (as depicted in [Figure 10](#) of Section 2.3).

The rest of the functional elements in [Figure 34](#) implement the Control Plane (CP). Of these, two represent the majority of the functionality that’s unique to the Mobile Core CP (as sketched in [Figure 13](#) of Section 2.4):

- *AMF (Access and Mobility Management Function)*: Responsible for connection and reachability management, mobility management, access authorization, and location services.
- *SMF (Session Management Function)*: Manages each UE session, including IP address allocation, selection of associated UP

function, control aspects of QoS, and control aspects of UP routing.

In other words, the AMF authorizes access when a UE first connects to one of the local base stations, and then tracks (but does not control) which base station currently serves each UE. The SMF then allocates an IP address to each AMF-authorized UE, and directly interacts with the UPF to maintain per-device session state.

Of particular note, the per-UE session state controlled by the SMF (and implemented by the UPF) includes a packet buffer in which packets destined to an idle UE are queued during the time the UE transitions to active state. This feature was originally designed to avoid data loss during a voice call, but its value is less obvious when the data is an IP packet since end-to-end protocols like TCP are prepared to retransmit lost packets. On the other hand, if idle-to-active transitions are too frequent, they can be problematic for TCP.

Before continuing with our inventory of control-related elements in [Figure 34](#), it is important to note we show only a fraction of the full set that 3GPP defines. The full set includes a wide range of possible features, many of which are either speculative (i.e., identify potential functionality) or overly prescriptive (i.e., identify well-known cloud native microservices). We limit our discussion to functional elements that provide value in the private 5G deployments that we focus on. Of these, several provide functionality similar to what one might find in any microservice-based application:

- *AUSF (Authentication Server Function)*: Authenticates UEs.
- *UDM (Unified Data Management)*: Manages user identity, including the generation of authentication credentials and access authorization.
- *UDR (Unified Data Repository)*: Manages user static subscriber-related information.
- *UDSF (Unstructured Data Storage Function)*: Used to store unstructured data, and so is similar to a *key-value store*.
- *NEF (Network Exposure Function)*: Exposes select capabilities to third-party services, and so is similar to an *API Server*.
- *NRF (Network Repository Function)*: Used to discover available

services (network functions), and so is similar to a *Discovery Service*.

The above list includes 3GPP-specified control functions that are, in some cases, similar to well-known microservices. In such cases, substituting an existing cloud native component is a viable implementation option. For example, MongoDB can be used to implement a UDSF. In other cases, however, such a one-for-one swap is not possible due to assumptions 3GPP makes. For example, AUSF, UDM, UDR, and AMF collectively implement a *Authentication and Authorization Service*, but an option like OAuth2 could not be used in their place because (a) UDM and UDR are assumed to be part of the global identity mapping service discussed in Section 5.1, and (b) 3GPP specifies the interface by which the various components request service from each other (e.g., AMF connects to the RAN via the N2 interface depicted in [Figure 34](#)). We will see how to cope with such issues in Section 5.3, where we talk about implementation issues in more detail.

Finally, [Figure 34](#) shows two other functional elements that export a northbound interface to the management plane (not shown):

- *PCF (Policy Control Function)*: Manages the policy rules for the rest of the Mobile Core CP.
- *NSSF (Network Slice Selection Function)*: Manages how network slices are selected to serve a given UE.

Keep in mind that even though 3GPP does not directly prescribe a microservice implementation, the overall design clearly points to a cloud native solution as the desired end-state for the Mobile Core. Of particular note, introducing a distinct storage service means that all the other services can be stateless, and hence, more readily scalable.

5.3 Control Plane

This section describes two different strategies for implementing the Mobile Core CP. Both correspond to open source projects that are available for download and experimentation.

5.3.1 SD-Core

Our first example, called SD-Core, is a nearly one-for-one translation of the functional blocks shown in [Figure 34](#) into a cloud native implementation. A high-level schematic is shown in [Figure 35](#), where each element corresponds to a scalable set of Kubernetes-hosted containers. We include this schematic even though it looks quite similar to [Figure 34](#) because it highlights four implementation details.

Further Reading

[SD-Core](#).

Figure 35. SD-Core implementation of the Mobile Core Control Plane, including support for Standalone (SA) deployment of both 4G and 5G.

First, SD-Core supports both the 5G and 4G versions of the Mobile Core, [\[1\]](#) which share a common User Plane (UPF). We have not discussed details of the 4G Core, but observe that it is less disaggregated. In particular, the components in the 5G Core are specified so that they can be stateless, simplifying the task of horizontally scaling them out as load dictates. (The rough correspondence between 4G and 5G is: MME-to-AMF, SPGW_C-to-SMF, HSS-to-UDM, and PCRF-to-PCF.) Although not shown in the schematic, there is also a scalable key-value store microservice based on MongoDB. It is used to make Core-related state persistent for the Control Planes; for example, UDM/UDR (5G) and HSS (4G) write subscriber state to MongoDB.

SD-Core's 4G Core is a fork of the OMEC project and its 5G Core is a fork of the Free5GC project.

Second, [Figure 35](#) illustrates 3GPP's *Standalone (SA)* deployment option, in which 4G and 5G networks co-exist and run independently. They share a UPF implementation, but UPF instances are instantiated separately for each RAN/Core pair, with support for both the 4G and 5G interfaces, denoted *S1-U* and *N3*, respectively. Although not obvious from the SA example, 3GPP

defines an alternative transition plan, called *NSA (Non-Standalone)*, in which separate 4G and 5G RANs were paired with either a 4G Core or a 5G Core. The details of how that works are not relevant to this discussion, except to make the point that production networks almost never get to enjoy a “flag day” on which a new version is universally substituted for an old version. A migration plan has to be part of the design. More information on this topic can be found in a GSMA Report.

Further Reading

[Road to 5G: Introduction and Migration](#). GSMA Report, April 2018.

Third, [Figure 35](#) shows many of the 3GPP-defined inter-component interfaces. These include an over-the-air interface between base stations and UEs (*NR Uu*), control interfaces between the Core and both UEs and base stations (*N1* and *N2*, respectively), a user plane interface between the Core and base stations (*N3*), and a data plane interface between the Core and the backbone network (*N6*).

The schematic also shows interfaces between the individual microservices that make up the Core’s Control Plane; for example, *Nudm* is the interface to the UDM microservice. These latter interfaces are RESTful, meaning clients access each microservice by issuing GET, PUT, POST, PATCH, and DELETE operations over HTTP, where a service-specific schema defines the available resources that can be accessed. Note that some of these interfaces are necessary for interoperability (e.g., *N1* and *N Uu* make it possible to connect your phone to any MNO’s network), but others could be seen as internal implementation details. We’ll see how Magma takes advantage of this distinction in the next section.

Fourth, by adopting a cloud native design, SD-Core benefits from being able to horizontally scale individual microservices. But realizing this benefit isn’t always straightforward. In particular, because the AMF is connected to the RAN by SCTP (corresponding to the *N1* and *N2* interfaces shown in [Figure 35](#)), it is necessary to put an *SCTP load balancer* in front of the AMF. This load balancer terminates the SCTP connections, and distributes requests across a set of AMF containers. These AMF instances then depend on a scalable backend store (specifically MongoDB) to read and write

shared state.

5.3.2 Magma

Magma is an open source Mobile Core implementation that takes a different and slightly non-standard approach. Magma is similar to SD-Core in that it is implemented as a set of microservices, but it differs in that it is designed to be particularly suitable for remote and rural environments with poor backhaul connectivity. This emphasis, in turn, leads Magma to (1) adopt an SDN-inspired approach to how it separates functionality into centralized and distributed components, and (2) factor the distributed functionality into microservices without strict adherence to all the standard 3GPP interface specifications. This refactoring is also a consequence of Magma being designed to unify 4G, 5G, and WiFi under a single architecture.

One of the first things to note about Magma is that it takes a different view of “backhaul” from the approaches we have seen to date. Whereas the backhaul networks shown previously connect the eNBs/gNBs and radio towers back to the mobile core ([Figure 3](#)), Magma actually puts much of the mobile core functionality right next to the radio as seen in [Figure 37](#). It is able to do this because of the way it splits the core into centralized and distributed parts. So Magma views “backhaul” as the link that connects a remote deployment to the rest of the Internet (including the central components), contrasting with conventional 3GPP usage. As explored further below, this can overcome many of the challenges that unreliable backhaul links introduce in conventional approaches.

Figure 36. Overall architecture of the Magma Mobile Core, including support for 4G and 5G, and Wi-Fi. There is one central Orchestrator and typically many Access Gateways (AGWs).

[Figure 36](#) shows the overall Magma architecture. The central part of Magma is the single box in the figure marked *Central Control & Management (Orchestrator)*. This is roughly analogous to the central controller found in typical SDN systems, and provides a northbound

API by which an operator or other software systems (such as a traditional OSS/BSS) can interact with the Magma core. The orchestrator communicates over backhaul links with Access Gateways (AGWs), which are the distributed components of Magma. A single AGW typically handles a small number of eNBs/gNBs. As an example, see [Figure 37](#) which includes a single eNB and AGW located on a radio tower. In this example, a point-to-point wireless link is used for backhaul.

The AGW is designed to have a small footprint, so that small deployments do not require a datacenter's worth of equipment. Each AGW also contains both data plane and control plane elements. This is a little different from the classic approach to SDN systems in which only the data plane is distributed. Magma can be described as a hierarchical SDN approach, as the control plane itself is divided into a centralized part (running in the Orchestrator) and a distributed part (running in the AGW). [Figure 36](#) shows the distributed control plane components and data plane in detail. We postpone a general discussion of orchestration until Chapter 6.

Figure 37. A sample Magma deployment in rural Peru, showing (a) point-to-point wireless backhaul, (b) LTE radio and antenna, (c) ruggedized embedded PC serving as AGW, and (d) solar power and battery backup for site.

Magma differs from the standard 3GPP approach in that it terminates 3GPP protocols logically close to the edge, which in this context corresponds to two interface points: (1) the radio interface connecting Magma to an eNB or gNB (implemented by set of modules on the left side of the AGW in the figure) or the federation interface connecting Magma to another mobile network (implemented by the *Federation Gateway* module in the figure). Everything “between” those two external interfaces is free to deviate from the 3GPP specification, which has a broad impact as discussed below.

One consequence of this approach is that Magma can interoperate with other implementations *only* at the edges. Thus, it is possible to connect a Magma mobile core to any standards-compliant 4G or 5G base station and expect it to work, and similarly, it is possible to federate a Magma core with an existing MNO's 4G or 5G network.

However, since Magma does not implement all the 3GPP interfaces that are internal to a mobile packet core, it is not possible to arbitrarily mix and match components within the core. Whereas (in principle) a traditional 3GPP implementation would permit an AMF from one vendor to interoperate with the SMF of another vendor, it is not possible to connect parts of a mobile core from another vendor (or another open source project) with parts of Magma, aside from via the two interfaces just described.

Being free to deviate from the 3GPP spec means Magma can take a unifying approach across multiple wireless technologies, including 4G, 5G and WiFi. There is a set of functions that the core must implement for any radio technology (e.g., finding the appropriate policy for a given subscriber by querying a database); Magma provides them in an access-technology-independent way. These functions form the heart of an Access Gateway (AGW), as illustrated on the right side of [Figure 36](#). On the other hand, control protocols that are specific to a given radio technology are terminated in technology-specific modules close to the radio. For example, *SCTP* shown on the left side of the figure is the RAN tunneling protocol introduced in Section 2.3. These technology-specific modules then communicate with the generic functions (e.g., subscriber management, access control and management) on the right using gRPC messages that are technology-agnostic.

Magma's design is particularly well suited for environments where backhaul links are unreliable, for example, when a satellite is used. This is because the 3GPP protocols that traditionally have to traverse the backhaul from core to eNB/gNB are quite sensitive to loss and latency. Loss or latency can cause connections to be dropped, which in turn forces UEs to repeat the process of attaching to the core. In practice, not all UEs handle this elegantly, sometimes ending up in a “stuck” state.

Magma addresses the challenge of unreliable backhaul in several ways. First, Magma frequently avoids sending messages over the backhaul entirely by running more functionality in the AGW, which is located close to the radio as seen above. Functions that would be centralized in a conventional 3GPP implementation are distributed out to the access gateways in Magma. Thus, for example, the

operations required to authenticate and attach a UE to the core can typically be completed using information cached locally in the AGW, without any traffic crossing the backhaul. Secondly, when Magma does need to pass information over a backhaul link (e.g., to obtain configuration state from the orchestrator), it does so using gRPC, which is designed to operate reliably in the face of unreliable or high-latency links.

Note that while Magma has distributed much of the control plane out to the AGWs, it still supports centralized management via the Orchestrator. For example, adding a new subscriber to the network is done centrally, and the relevant AGW then obtains the necessary state to authenticate that subscriber when their UE tries to attach to the network.

Finally, Magma adopts a *desired state* model for managing runtime and configuration state. By this we mean that it communicates a state change (e.g., the addition of a new session in the user plane) by specifying the desired end state via an API call. This is in contrast with the *incremental update* model that is common in the 3GPP specification. When the desired end state is communicated, the loss of a message or failure of a component has less serious consequences. This makes reasoning about changes across elements of the system more robust in the case of partial failures, which are common in challenged environments like the ones Magma is designed to serve.

Consider an example where we are establishing user plane state for a set of active sessions. Initially, there are two active sessions, X and Y. Then a third UE becomes active and a session Z needs to be established. In the incremental update model, the control plane would instruct the user plane to “add session Z”. The desired state model, by contrast, communicates the entire new state: “the set of sessions is now X, Y, Z”. The incremental model is brittle in the face of failures. If a message is lost, or a component is temporarily unable to receive updates, the receiver falls out of sync with the sender. So it is possible that the control plane believes that sessions X, Y and Z have been established, while the user plane has state for only X and Y. By sending the entire desired end state, Magma ensures that the receiver comes back into sync with the sender once

it is able to receive messages again.

As described, this approach might appear inefficient because it implies sending complete state information rather than incremental updates. However, at the scale of an AGW, which handles on the order of hundreds to a few thousands of subscribers, it is possible to encode the state efficiently enough to overcome this drawback. With the benefit of experience, mechanisms have been added to Magma to avoid overloading the orchestrator, which has state related to all subscribers in the network.

The desired state approach is hardly novel but differs from typical 3GPP systems. It allows Magma to tolerate occasional communication failures or component outages due to software restarts, hardware failures, and so on. Limiting the scope of 3GPP protocols to the very edge of the network is what enables Magma to rethink the state synchronization model. The team that worked on Magma describes their approach in more detail in an NSDI paper.

Further Reading

S. Hasan, *et al.* [Building Flexible, Low-Cost Wireless Access Networks With Magma](#). NSDI, April 2023.

Finally, while we have focused on its Control Plane, Magma also includes a User Plane component. The implementation is fairly simple, and is based on Open vSwitch (OVS). Having a programmable user plane is important, as it needs to support a range of access technologies, and at the same time, OVS meets the performance needs of AGWs. However, this choice of user plane is not fundamental to Magma, and other implementations have been considered. We take a closer look at the User Plane in the next section.

5.4 User Plane

The User Plane of the Mobile Core—corresponding to the UPF component in [Figure 34](#)—connects the RAN to the Internet. Much like the data plane for any router, the UPF forwards IP packets, but because UEs often sleep to save power and may be in the process of

being handed off from one base station to another, it sometimes has to buffer packets for an indeterminate amount of time. Also like other routers, a straightforward way to understand the UPF is to think of it as implementing a collection of Match/Action rules, where the UPF first classifies each packet against a set of matching rules, and then executes the associated action.

Using 3GPP terminology, packet classification is defined by a set of *Packet Detection Rules (PDRs)*, where a given PDR might simply match the device's IP address, but may also take the domain name of the far end-point into consideration. Each attached UE has at least two PDRs, one for uplink traffic and one for downlink traffic, plus possibly additional PDRs to support multiple traffic classes (e.g., for different QoS levels, pricing plans, and so on.). The Control Plane creates, updates, and removes PDRs as UEs attach, move, and detach.

Each PDR then identifies one or more actions to execute, which in 3GPP terminology are also called “rules”, of which there are four types:

- **Forwarding Action Rules (FARs):** Instructs the UPF to forward downlink packets to a particular base station and uplink traffic to a next-hop router. Each FAR specifies a set of parameters needed to forward the packet (e.g., how to tunnel downlink packets to the appropriate base station), plus one of the following processing flags: a *forward* flag indicates the packet should be forwarded up to the Internet; a *tunnel* flag indicates the packet should be tunneled down to a base station; a *buffer* flag indicates the packet should be buffered until the UE becomes active; and a *notify* flag indicates that the CP should be notified to awaken an idle UE. FARs are created and removed when a device attaches or detaches, respectively, and the downlink FAR changes the processing flag when the device moves, goes idle, or awakes.
- **Buffering Action Rules (BARs):** Instructs the UPF to buffer downlink traffic for idle UEs, while also sending a *Downlink Data Notification* to the Control Plane. This notification, in turn, causes the CP to instruct the base station to awaken the UE. Once the UE becomes active, the UPF releases the

buffered traffic and resumes normal forwarding. The buffering and notification functions are activated by modifying a FAR to include *buffer* and *notify* flags, as just described. An additional set of parameters are used to configure the buffer, for example setting its maximum size (number of bytes) and duration (amount of time). Optionally, the CP can itself buffer packets by creating a PDR that directs the UPF to forward data packets to the control plane.

- **Usage Reporting Rules (URRs):** Instructs the UPF to periodically send usage reports for each UE to the CP. These reports include counts of the packets sent/received for uplink/downlink traffic for each UE and traffic class. These reports are used to both limit and bill subscribers. The CP creates and removes URRs when the device attaches and detaches, respectively, and each URR specifies whether usage reports should be sent periodically or when a quota is exceeded. A UE typically has two URRs (for uplink/downlink usage), but if a subscriber's plan includes special treatment for certain types of traffic, an additional URR is created for each such traffic class.
- **Quality Enforcement Rules (QERs):** Instructs the UPF to guarantee a minimum amount of bandwidth and to enforce a bandwidth cap. These parameters are specified on a per-UE / per-direction / per-class basis. The CP creates and removes QERs when a device attaches and detaches, respectively, and modifies them according to operator-defined events, such as when the network becomes more or less congested, the UE exceeds a quota, or the network policy changes (e.g., the user signs up for a new pricing plan). The UPF then performs traffic policing to enforce the bandwidth cap, along with packet scheduling to ensure a minimum bandwidth in conjunction with admission control in the control plane.

The rest of this section describes two complementary strategies for implementing a UPF, one server-based and one switch-based.

5.4.1 *Microservice Implementation*

A seemingly straightforward approach to supporting the set of

Match/Action rules just described is to implement the UPF in software on a commodity server. Like any software-based router, the process would read a packet from an input port, classify the packet by matching it against a table of configured PDRs, execute the associated action(s), and then write the packet to an output port. Such a process could then be packaged as a Docker container, with one or more instances spun up on a Kubernetes cluster as workload dictates. This is mostly consistent with a microservice-based approach, with one important catch: the actions required to process each packet are stateful.

What we mean by this is that the UPF has two pieces of state that needs to be maintained on a per-UE / per-direction / per-class basis: (1) a finite state machine that transitions between *forward*, *tunnel*, *buffer*, and *notify*; and (2) a corresponding packet buffer when in *buffer* state. This means that as the UPF scales up to handle more and more traffic—by adding a second, third, and fourth instance—packets still need to be directed to the original instance that knows the state for that particular flow. This breaks a fundamental assumption of a truly horizontally scalable service, in which traffic can be randomly directed to any instance in a way that balances the load. It also forces you to do packet classification before selecting which instance is the right one, which can potentially become a performance bottleneck, although it is possible to offload the classification stage to a SmartNIC/IPU.

5.4.2 P4 Implementation

Since the UPF is fundamentally an IP packet forwarding engine, it can also be implemented—at least in part—as a P4 program running on a programmable switch. Robert MacDavid and colleagues describe how that is done in SD-Core, which builds on the base packet forwarding machinery described in our companion SDN book. For the purposes of this section, the focus is on the four main challenges that are unique to implementing the UPF in P4.

Further Reading

R. MacDavid, *et al.* [A P4-based 5G User Plane Function](#). Symposium

on SDN Research, September 2021.

[Software-Defined Networks: A Systems Approach](#). November 2021.

First, P4-programmable forwarding pipelines include an explicit “matching” mechanism built on *Ternary Content-Addressable Memory* (TCAM). This memory supports fast table lookups for patterns that include wildcards, making it ideal for matching IP prefixes. In the case of the UPF, however, the most common PDRs correspond to exact matches of IP addresses (for downlink traffic to each UE) and GTP tunnel identifiers (for uplink traffic from each UE). More complex PDRs might include regular expressions for DNS names or require deep packet inspection.

Because TCAM capacity is limited, and the number of unique PDRs that need to be matched in both directions is potentially in the tens of thousands, it’s necessary to use the TCAM judiciously. One implementation strategy is to set up two parallel PDR tables: one using the relatively plentiful switch SRAM for common-case uplink rules that exactly matches on tunnel identifiers (which can be treated as table indices); and one using TCAM for common-case downlink rules that match the IP destination address.

Second, when a packet arrives from the Internet destined for an idle UE, the UPF buffers the packet and sends an alert to the 5G control plane, asking that the UE be awakened. Today’s P4-capable switches do not have large buffers or the ability to hold packets indefinitely, but a buffering microservice running on a server can be used to address this limitation. The microservice indefinitely holds any packets that it receives, and later releases them back to the switch when instructed to do so. The following elaborates on how this would work.

When the Mobile Core detects that a UE has gone idle (or is in the middle of a handover), it creates a FAR with the *buffer* flag set, causing the on-switch P4 program to redirect packets to the buffering microservice. Packets are redirected without modifying their IP headers by placing them in a tunnel, using the same tunneling protocol that is used to send data to base stations. This allows the switch to treat the buffering microservice just like another base station.

When the first packet of a flow arrives at the buffering microservice, it sends an alert to the CP, which then (1) wakes up the UE, (2) modifies the corresponding FAR by unsetting the *buffer* flag and setting the *tunnel* flag, and once the UE is active, (3) instructs the buffering microservice to release all packets back to the switch. Packets arriving at the switch from the buffering microservice skip the portion of the UPF module they encountered before buffering, giving the illusion they are being buffered in the middle of the switch. That is, their processing resumes at the tunneling stage, where they are encapsulated and routed to the appropriate base station.

Third, QERs cannot be fully implemented in the switch because P4 does not include support for programming the packet scheduler. However, today's P4 hardware does include fixed-function schedulers with configurable weights and priorities; these parameters are set using a runtime interface unrelated to P4. A viable approach, similar to the one MacDavid, Chen, and Rexford describe in their INFOCOM paper, is to map each QoS class specified in a QER onto one of the available queues, and assign a weight to that queue proportional to the fraction of the available bandwidth the class is to receive. As long as each class/queue is not over subscribed, individual UEs in the class will receive approximately the bit rate they have been promised. As an aside, since 3GPP under-specifies QoS guarantees (leaving the details to the implementation), such an approach is 3GPP-compliant.

Further Reading


R. MacDavid, X. Chen, J. Rexford. [Scalable Real-time Bandwidth Fairness in Switches](#). IEEE INFOCOM, May 2023.


Finally, while the above description implies the Mobile Core's CP talks directly to the P4 program on the switch, the implementation is not that straightforward. From the Core's perspective, the SMF is responsible for sending/receiving control information to/from the UPF, but the P4 program implementing the UPF is controlled through an interface (known as P4Runtime or P4RT) that is auto-generated from the P4 program being controlled. MacDavid's paper describes how this is done in more detail (and presumes a deep understanding of the P4 toolchain), but it can be summarized as


follows. It is necessary to first write a “Model UPF” in P4, use that to program to generate the UPF-specific P4RT interface, and then write translators that (1) connect SMF to P4RT, and (2) connect P4RT to the underlying physical switches and servers. A high-level schematic of this software stack is shown in [Figure 38](#).

Figure 38. A model P4-based implementation of the UPF is used to generate the interface that is then used by the SMF running in the Mobile Core control plane to control the physical implementation of the UPF running on a combination of hardware switches and servers.


Note that while this summary focuses on how the CP controls the UPF (the downward part of the schematic shown in [Figure 38](#)), the usage counters needed to generate URRs that flow upward to the CP are easy to support because the counters implemented in the switching hardware are identical to the counters in the Model UPF. When the Mobile Core requests counter values from the Model UPF, the backend translator polls the corresponding hardware switch counters and relays the response.







Manage end-to-end connectivity on a per-device/application basis.



Chapter 6: Managed

Cloud Service

This chapter describes how to assemble all the pieces described in the previous chapters to provide 5G connectivity as a managed cloud service. Such a service might be deployed in enterprises, for example, in support of collection of operational data, video, robots, IoT devices, and so on—a set of use cases sometimes referred to as Industry 4.0.

The first step is to implement all the components using cloud native building blocks. We start by introducing those building blocks in Section 6.1. The second step is to introduce yet another component—a *Cloud Management Platform*—that is responsible for operationalizing 5G-as-a-Service. The rest of the sections describe how to build such a management system using open source tools.

Before getting into the details, it is important to remember that mobile cellular service (both voice and broadband) has been offered as a Telco service for 40 years. Treating it as a managed cloud service is a significant departure from that history, most notably in how the connectivity it provides is operated and managed. As a consequence, the Cloud Management Platform described in this chapter is significantly different from the legacy OSS/BSS mechanisms that have traditionally been the centerpiece of the Telco management machinery. The terminology is also different, but that only matters if you are trying to map Telco terminology onto cloud terminology (which we are not). We take up the “terminology mapping problem” in a companion book, and here focus instead on a from-scratch cloud-based design.

Further Reading

L. Peterson, A. Bavier, S. Baker, Z. Williams, and B. Davie. [Edge Cloud Operations: A Systems Approach](#). June 2022.

6.1 Building Blocks

The implementation strategy starts with commodity hardware and open source software. These building blocks will be familiar to anyone who has built a cloud native application, but they deserve to be explicitly named in a discussion of mobile cellular networks, which have historically been built using closed, proprietary hardware devices.

The hardware building blocks include bare-metal servers and switches, which might include ARM or x86 processor chips and Tomahawk or Tofino switching chips, respectively. A physical cloud cluster is then constructed with the hardware building blocks arranged as shown in [Figure 39](#): one or more racks of servers connected by a leaf-spine switching fabric. We show the servers above the switching fabric to emphasize that software running on the servers controls the switches (as we will see in the next section).

Figure 39. Example building block components used to construct an edge cloud, including commodity servers and switches, interconnected by a leaf-spine switching fabric.

The software building blocks start with the following open source components:

1. Docker containers package software functionality.
2. Kubernetes instantiates and interconnects a set of containers.
3. Helm specifies how collections of related containers are interconnected to build microservice-based applications.
4. Fleet specifies how a set of Kubernetes applications are to be deployed on the available infrastructure.
5. Terraform provisions a set of one or more Kubernetes clusters, configuring them to host microservice applications.

Docker is a container runtime that leverages OS isolation APIs to instantiate and run multiple containers, each of which is an instance defined by a Docker image. Docker images are most frequently built using a Dockerfile, which uses a layering approach that allows sharing and building customized images on top of base images. A final image for a particular task incorporates all dependencies required by the software that is to run in the container, resulting in a container image that is portable across servers, depending only on the kernel and Docker runtime. We also assume one or more image

artifact repositories of Docker containers that we will want to deploy in our cloud, of which <https://hub.docker.com/> is the best known example.

Further Reading

[Docker Tutorial](#).

Kubernetes is a container orchestration system. It provides a programmatic interface for scaling container instances up and down, allocating server resources to them, setting up virtual networks to interconnect those instances, and opening service ports that external clients can use to access those instances. Behind the scenes, Kubernetes monitors the liveness of those containers, and automatically restarts any that have failed. In other words, if you instruct Kubernetes to spin up three instances of microservice X, Kubernetes will do its best to keep three instances of the container that implements X running at all times.

Further Reading

[Kubernetes Tutorial](#).

Helm is a configuration manager that runs on top of Kubernetes. It issues calls against the Kubernetes API according to a developer-provided specification, known as a *Helm Chart*. It is now common practice for cloud applications built from a set of microservices to publish a Helm chart that defines how the application is to be deployed on a Kubernetes cluster. See <https://artifacthub.io/> for a collection of publicly available Helm Charts.

Further Reading

[Helm Tutorial](#).

Fleet, an application deployment manager, is responsible for installing a *Bundle* of Helm Charts on one or more target clusters. If we were trying to deploy a single Chart on just one Kubernetes cluster, then Helm would be sufficient. The value of Fleet is that it scales up that process, helping us manage the deployment of multiple charts across multiple clusters. Moreover, Fleet does this

using an approach known as *Configuration-as-Code*, where the desired configuration is checked into a repo, just like any other software. Checking a new or updated Bundle into a repo triggers the deployment of the corresponding applications.

Further Reading

[Fleet: GitOps at Scale.](#)

Terraform is an infrastructure manager that, in our scenario, provisions one or more Kubernetes clusters, preparing them to host a collection of Helm-specified applications. It does this using an approach known as *Infrastructure-as-Code*, which documents exactly how the infrastructure is to be configured in a declarative format that can be (a) checked into a repo, (b) version-controlled, and (c) executed just like any piece of software. Terraform assumes an underlying provisioning API, with Microsoft's Azure Kubernetes Service (AKS), AWS's Amazon Elastic Kubernetes Service (EKS), Google's Google Kubernetes Engine (GKE) and Rancher's Rancher Kubernetes Engine (RKE) being widely available examples.

Further Reading

[Terraform Tutorials.](#)

The inter-related roles of Helm, Fleet, and Terraform can be confusing, in part because there is overlap in what each tries to do. One distinction is that Helm Charts are typically specified by *developers* as a way of specifying how an application is constructed from a set of microservices, whereas Fleet and Terraform give *operators* an opportunity to specify details of their particular deployment scenarios. A second distinction is that Helm and Fleet help manage the *applications running on* one or more Kubernetes clusters, whereas Terraform is used to set up and configure the *underlying Kubernetes clusters* in the first place. Again, there is overlap in the capabilities of these respective tools, but these two distinctions characterize how they are used in Aether. The more general takeaway is that cloud management has to accommodate both developers and operators, and to clearly delineate between applications and platforms.

6.2 Example Deployment

Using these building blocks, it is possible to construct a wide range of deployment scenarios for a managed 5G service. For illustrative purposes, we use a particular deployment based on the Aether edge cloud introduced in Chapter 2. Aether is an operational edge cloud that has been deployed to multiple sites, and most importantly for our purposes, includes an API that edge apps can use to customize 5G connectivity to better meet their objectives.

6.2.1 Edge Cloud

An Aether edge deployment, called ACE (Aether Connected Edge), is a Kubernetes-based cluster. It consists of one or more server racks interconnected by a leaf-spine switching fabric, with an SDN control plane (denoted SD-Fabric) managing the fabric. We briefly saw SD-Fabric in Chapter 5 as an implementation option for the Mobile Core's User Plane Function (UPF), but for an in-depth description of SD-Fabric, we refer you to a companion book.

Further Reading

L. Peterson, C. Cascone, B. O'Connor, T. Vachuska, and B. Davie. [Software-Defined Networks: A Systems Approach](#). November 2021.

Figure 40. Aether Connected Edge (ACE) = The cloud platform (Kubernetes and SD-Fabric) plus the 5G connectivity service (RAN and User Plane of Mobile Core). Dotted lines (e.g., between SD-RAN and the individual base stations, and between the Network OS and the individual switches) represent control relationships (e.g., SD-RAN controls the small cells and SD-Fabric controls the switches).

As shown in [Figure 40](#), ACE hosts two additional microservice-based subsystems on top of this platform; they collectively implement *5G-as-a-Service*. The first subsystem, SD-RAN, is the SDN-based implementation of the Radio Access Network described in Chapter 4. It controls the small cell base stations deployed throughout the enterprise. The second subsystem, SD-Core, is an

SDN-based implementation of the User Plane half of the Mobile Core described in Chapter 5. It is responsible for forwarding traffic between the RAN and the Internet. The SD-Core Control Plane (CP) runs off-site, and is not shown in [Figure 40](#). Both subsystems (as well as the SD-Fabric), are deployed as a set of microservices, just as any other cloud native workload.

Once an edge cluster is running in this configuration, it is ready to host a collection of cloud-native edge applications (not shown in [Figure 40](#)). What's unique to our example configuration is its ability to connect such applications to mobile devices throughout the enterprise using the 5G Connectivity Service implemented by SD-RAN and SD-Core, without the resulting network traffic ever leaving the enterprise; a scenario known as *local breakout*. Moreover, this service is offered as a managed service, with enterprise system administrators able to use a programmatic API (and associated GUI portal) to control that service; that is, authorize devices, restrict access, set QoS profiles for different devices and applications, and so on.

6.2.2 Hybrid Cloud

While it is possible to instantiate a single ACE cluster in just one site, Aether is designed to support multiple edge deployments, all of which are managed from the central cloud. Such a hybrid cloud scenario is depicted in [Figure 41](#), which shows two subsystems running in the central cloud: (1) one or more instantiations of the Mobile Core Control Plane (CP), and (2) the Aether Management Platform (AMP).

Each SD-Core CP controls one or more SD-Core UPFs. Exactly how CP instances (running centrally) are paired with UPF instances (running at the edges) is a runtime decision, and depends on the degree of isolation the enterprise sites require. AMP is Aether's realization of a Cloud Management Platform; it is responsible for managing all the centralized and edge subsystems (as introduced in the next section).

Figure 41. Aether runs in a hybrid cloud configuration, with Control

Plane of Mobile Core and the Aether Management Platform (AMP) running in the Central Cloud.

There is an important aspect of this hybrid cloud that is not obvious from [Figure 41](#), which is that the “hybrid cloud” we keep referring to is best described as a set of Kubernetes clusters, rather than a set of physical clusters. This is because, while each ACE site usually corresponds to a physical cluster built out of bare-metal components, each of the SD-Core CP subsystems shown in [Figure 41](#) is actually deployed in a logical Kubernetes cluster on a commodity cloud. The same is true for AMP. Aether’s centralized components are able to run in Google Cloud Platform, Microsoft Azure, and Amazon’s AWS. They also run as an emulated cluster implemented by a system like KIND (Kubernetes in Docker), making it possible for developers to run these components on their laptops.

Near-Edge vs Far-Edge

We use enterprises as the exemplar edge deployment in this book, without prescribing a role for traditional MNOs. When traditional MNOs are involved, it is not uncommon for them to make a distinction between the “near-edge” and the “far-edge”, where the far-edge corresponds to the enterprise and the near-edge corresponds to their traditional aggregation points (or Central Offices), as described in Section 1.2. In such a scenario, it is typically the case that the RU and DU are located at the far-edge (on-prem), while the CU—along with both the Control and User Planes of the Mobile Core—run in the near-edge. Such a configuration does not support local breakout, since all traffic must travel to the near-edge before being routed to the edge app (which might be running back in the enterprise).

In contrast, the deployment described in this Chapter has everything except the Mobile Core Control Plane (CP) running on-prem. Moreover, because there is no traditional MNO involved, there is no near-edge to speak of, with the Core CP instead running in a central cloud. For example, this section describes a deployment with SD-Core (CP) running in the Google Cloud. It is the case, however, that the SD-Core (CP) can optionally run on-prem if a fully local configuration is preferred. Where each component runs is a configuration option.

6.2.3 Stakeholders

With the understanding that our target environment is a collection of Kubernetes clusters—some running on bare-metal hardware at edge sites and some running in central datacenters—there is an orthogonal issue of how decision-making responsibility for those clusters is shared among multiple stakeholders. Identifying the relevant stakeholders is an important prerequisite for establishing a cloud service, and while the example we use may not be suitable for all situations, it does illustrate the design implications.

For Aether, we care about two primary stakeholders: (1) the *cloud operators* who manage the hybrid cloud as a whole, and (2) the *enterprise users* who decide on a per-site basis how to take advantage of the local cloud resources (e.g., what edge applications to run and how to slice connectivity resources among those apps). We sometimes call the latter “enterprise admins” to distinguish them from “end-users” who might want to manage their own personal devices.

Aether is multi-tenant in the sense that it authenticates and isolates these stakeholders, allowing each to access only those objects they are responsible for. This makes the approach agnostic as to whether all the edge sites belong to a single organization (with that organization also responsible for operating the cloud), or alternatively, there being a separate organization that offers a managed service to a set of distinct enterprises (each of which spans one or more sites).

There is a third stakeholder of note—third-party service providers—which points to the larger issue of how we deploy and manage the edge applications that take advantage of 5G-as-a-Service. The approach Aether adopts is to expect service providers to make their applications available either as source code (which works for open source or in-house apps), or as standard cloud native artifacts (e.g., Docker images and Helm charts). Either format can be fed into the Lifecycle Management pipeline described in Section 6.3.2. The alternative would be for edge service providers to share operational responsibility for the edge cloud with the cloud operator, which is possible if the infrastructure running at the edge is either multi-

tenant or a multi-cloud.

6.2.4 Alternative Configurations

The deployment just described is Aether in its full glory. Simpler configurations are also possible, which makes sense in less demanding scenarios. Examples include:

- Small edge clusters can be built with only a single switch (or two switches for resiliency), with or without SDN-based control. In the limit, an Aether edge can run on a single server.
- It is possible to substitute legacy small cells for O-RAN compliant small cells and the SD-RAN solution that includes a near RT-RIC and associated xApps.
- It is possible co-locate both AMP and the SD-Core on the edge cluster, resulting in a complete Aether deployment that is self-contained in a single site.

These are all straightforward configuration options. A very different approach is to start with an edge cluster that is managed by one of the hyperscalers, rather than have Aether provision Kubernetes on bare-metal. Google's Anthos, Microsoft's Azure Arc, and Amazon's ECS-Anywhere are examples of such edge cloud products. In such a scenario, AMP still manages the SD-Core and SD-RAN applications running on top of Kubernetes, but not the underlying platform (which may or may not include an SDN-based switching fabric).

Another variable in how 5G can be deployed at the edge is related to who owns the underlying cloud infrastructure. Instead of a cloud provider, an enterprise, or a traditional MNO owning the hardware, there are situations where a third-party, often called a *neutral host*, owns and operates the hardware (along with the real estate it sits in), and then rents access to these resources to multiple 5G providers. Each mobile service provider is then a tenant of that shared infrastructure.

This kind of arrangement has existed for years, albeit with conventional RAN devices, but shifting to a cloud-based design

makes it possible for neutral hosts to lease access to *virtualized* edge resources to their tenants. In principle, the only difference between this scenario and today's multi-tenant clouds is that such providers would offer edge resources—located in cell towers, apartment buildings, and dense urban centers—instead of datacenter resources. The business arrangements would also have to be different from Private 5G, but the technical design outlined in this book still applies.

6.3 Cloud Management Platform

Operationalizing the hardware and software components described in the previous two sections is the essence of what it means to offer 5G as a *managed service*. This responsibility falls to the Cloud Management Platform, which in Aether corresponds to the centralized AMP component shown in [Figure 41](#). AMP manages both the distributed set of ACE clusters and one or more SD-Core CP clusters running in the central cloud.

The following uses AMP to illustrate how to deliver 5G-as-a-Service, but the approach generalizes because AMP is based on widely-used open source tools. For more details about all the subsystems involved in operationalizing an edge cloud, we refer you to the companion book mentioned in the introduction to this chapter.

Figure 42. The four subsystems that comprise AMP: Resource Provisioning, Lifecycle Management, Service Orchestrator, and Monitoring & Telemetry.

At a high level, AMP is organized around the four subsystems shown in [Figure 42](#):

- **Resource Provisioning** is responsible for initializing resources (e.g., servers, switches) that add, replace, or upgrade capacity. It configures and bootstraps both physical and virtual resources, bringing them up to a state so Lifecycle Management can take over and manage the software running

on those resources.

- **Lifecycle Management** is responsible for continuous integration and deployment of the software components that collectively implement 5G-as-a-Service. It adopts the GitOps practice of *Configuration-as-Code*, using Helm Charts, Terraform Templates, and Fleet Bundles to specify how functionality is to be deployed and configured.
- **Service Orchestration** provides a means to manage services once they are operational. It defines an API that hides the implementation details of the underlying microservices, and is used to manage the provided 5G connectivity service.
- **Monitoring & Telemetry** is responsible for collecting, archiving, evaluating, and analyzing operational data generated by the underlying components. It makes it possible to diagnose and respond to failures, tune performance, do root cause analysis, perform security audits, and understand when it is necessary to provision additional capacity.

AMP implements all four subsystems, but an alternative perspective that characterizes the management platform as having *online* and *offline* components is also instructive. Such a two dimensional schematic is shown in [Figure 43](#). Lifecycle Management (coupled with Resource Provisioning) runs offline, sitting adjacent to the hybrid cloud. Operators and Developers provision and change the system by checking code (including configuration specs) into a repo, which in turn triggers an upgrade of the running system. Service Orchestration (coupled with Monitoring and Telemetry) runs online, layered on top of the hybrid cloud being managed. It defines an API that can be used to read and write parameters of the running system, which serves as a foundation for building closed-loop control.

Figure 43. Alternative representation of the management platform, highlighting the offline and online aspects of cloud management.

The offline and online aspects of cloud management are related in the sense that the offline component is also responsible for lifecycle-managing the online component. This is because the latter is deployed as a collection of Kubernetes applications, just like SD-Core and SD-RAN. Version management is a key aspect of this

relationship since the runtime API to the 5G connectivity service has to stay in sync with the underlying implementation of the constituent subsystems. How Aether realizes version control is described in more detail in the companion Edge Cloud Operations book.

6.3.1 Resource Provisioning

Resource Provisioning is the process of bringing virtual and physical resources online. For physical resources, it has both a hands-on component (racking and connecting devices) and a bootstrap component (configuring how the resources boot into a “ready” state). When utilizing virtual resources (e.g., VMs instantiated on a commercial cloud) the “rack and connect” step is carried out by a sequence of API calls rather than a hands-on technician.

Because we want to automate the sequence of calls needed to activate virtual infrastructure, we adopt an approach known as *Infrastructure-as-Code*. This is where Terraform comes into play. The general idea is to document, in a declarative format that can be “executed”, exactly what our infrastructure is to look like. The code defines how the infrastructure is to be configured.

When a cloud is built from a combination of virtual and physical resources, as is the case for a hybrid cloud like Aether, we need a seamless way to accommodate both. To this end, our approach is to first overlay a *logical structure* on top of hardware resources, making them roughly equivalent to the virtual resources we get from a commercial cloud provider. This results in a hybrid scenario similar to the one shown in [Figure 44](#). One way to think about this is that the task of booting hardware into the “ready” state involves installing and configuring several subsystems that collectively form the cloud platform. It is this platform that Terraform interacts with, indirectly, through a cloud provisioning API.

Figure 44. Resource Provisioning in a hybrid cloud that includes both physical and virtual resources.

6.3.2 Lifecycle Management

Lifecycle Management is concerned with updating and evolving a running system over time. [Figure 45](#) gives an overview of the pipeline/toolchain that make up the two halves of Lifecycle Management—Continuous Integration (CI) and Continuous Deployment (CD). The key thing to focus on is the Image and Config Repos in the middle. They represent the “interface” between the two halves: CI produces Docker Images and Helm Charts, storing them in the respective Repositories, while CD consumes Docker Images and Helm Charts, pulling them from the respective Repositories.

Figure 45. Overview of the CI/CD pipeline.

The Config Repo also contains declarative specifications of the infrastructure artifacts (specifically, Terraform templates and Fleet Bundles). These files are input to Lifecycle Management, which implies that Terraform and Fleet gets invoked as part of CI/CD whenever these files change. In other words, CI/CD keeps both the software-related components in the underlying cloud platform and the microservice workloads that run on top of that platform up to date.

Continuous Delivery vs Deployment

You will also hear CD refer to “Continuous Delivery” instead of “Continuous Deployment”, but we are interested in the complete end-to-end process, so CD will always imply the latter in this book. But keep in mind that “continuous” does not necessarily mean “instantaneous”; there can be a variety of gating functions injected into the CI/CD pipeline to control when and how upgrades get rolled out. The important point is that all the stages in the pipeline are automated.

So what exactly does “Continuous Delivery” mean? Arguably, it’s redundant when coupled with “Continuous Integration” since the set of artifacts being produced by the CI half of the pipeline (e.g., Docker images) is precisely what’s being delivered. There is no “next step” unless you also deploy those artifacts. It’s hair-splitting, but some would argue CI is limited to testing new code and Continuous Delivery corresponds to

the final “publish the artifact” step. For our purposes, we lump “publish the artifact” into the CI half of the pipeline.

There are three takeaways from this overview. The first is that by having well-defined artifacts passed between CI and CD (and between operators responsible for resource provisioning and CD), the subsystems are loosely coupled, and able to perform their respective tasks independently. The second is that all authoritative state needed to successfully build and deploy the system is contained within the pipeline, specifically, as declarative specifications in the Config Repo. This is the cornerstone of *Configuration-as-Code* (also known as *GitOps*), the cloud native approach to CI/CD. The third is that there is an opportunity for operators to apply discretion to the pipeline, as denoted by the “*Deployment Gate*” in the Figure, controlling what features get deployed when.

The third repository shown in [Figure 45](#) is the Code Repo (on the far left). Developers continually check new features and bug fixes into this repo, which triggers the CI/CD pipeline. A set of tests and code reviews are run against these check-ins, with the output of those tests/reviews reported back to developers, who modify their patch sets accordingly. (These develop-and-test feedback loops are implied by the dotted lines in [Figure 45](#).)

The far right of [Figure 45](#) shows the set of deployment targets, with *Staging* and *Production* called out as two illustrative examples. The idea is that a new version of the software is deployed first to a set of Staging clusters, where it is subjected to realistic workloads for a period of time, and then rolled out to the Production clusters once the Staging deployments give us confidence that the upgrade is reliable.

Finally, two of the CI stages shown in [Figure 45](#) identify a *Testing* component. One is a set of component-level tests that are run against each patch set checked into the Code Repo. These tests gate integration; fully merging a patch into the Code Repo requires first passing this preliminary round of tests. Once merged, the pipeline runs a build across all the components, and a second round of testing happens on a *Quality Assurance (QA)* cluster. Passing these tests gate deployment, but as just noted, testing also happens in the

Staging clusters as part of the CD end of the pipeline.

6.3.3 Service Orchestration

Service Orchestration is responsible for managing the Kubernetes workloads once they are up and running, which in our case means providing a programmatic API that can be used by various stakeholders to manage the 5G connectivity service. As shown in [Figure 46](#), the Service Orchestrator hides the implementation details of 5G connectivity, which spans four different components and multiple clouds. It does this by providing a coherent service interface for users, enabling them to authorize devices and set QoS parameters on an end-to-end basis.

Figure 46. Example use case that requires ongoing runtime control.

In other words, the Service Orchestrator defines an abstraction layer on top of a collection of backend components, effectively turning them into an externally visible (and controllable) cloud service. In some situations a single backend component might implement the entirety of a service, but in the case of 5G, which is constructed from a collection of disaggregated components, Service Orchestration is where we define an API that logically integrates those components into a unified and coherent whole. It is also an opportunity to “raise the level of abstraction” for the underlying subsystems, hiding unnecessary implementation details.

We describe this connectivity interface in Section 6.4. For now, our focus is on the main issues Service Orchestration must address in order to offer such an API. At a high level, it must:

1. Authenticate the principal wanting to perform the operation.
2. Determine if that principal has sufficient privilege to carry out the operation.
3. Push the new parameter setting(s) to one or more backend components.
4. Record the specified parameter setting(s), so the new value(s) persist.

Central to this role is the requirement that Service Orchestration be able to represent a set of abstract objects, which is to say, it implements a *data model*. The API is then generated from this data model, and persistent state associated with instances of the models is stored in a key-value store. Aether uses YANG to specify the models, in part because it is a rich language for data modeling, but also because there is a robust collection of YANG-based tools that we can build upon.

Further Reading

[YANG - A Data Modeling Language for the Network Configuration Protocol](#). RFC 6020. October 2010.

Finally, changes to the model-defined parameters must be propagated to the backend components, and in practice there is no established API for doing this. Aether assumes gNMI as its southbound interface to communicate configuration changes to the software services, where an Adapter (not shown in the figure) has to be written for any services that do not support gNMI natively.

6.3.4 Monitoring and Telemetry

Collecting telemetry data for a running system is an essential function of the management platform. It enables operators to monitor system behavior, evaluate performance, make informed provisioning decisions, respond to failures, identify attacks, and diagnose problems. There are three types of telemetry data—*metrics*, *logs*, and *traces*—along with open source software stacks available to help collect, store, and act upon each of them.

Metrics are quantitative data about a system. These include common performance metrics such as link bandwidth, CPU utilization, and memory usage, but also binary results corresponding to “up” and “down”, as well as other state variables that can be encoded numerically. These values are produced and collected periodically (e.g., every few seconds), either by reading a counter, or by executing a runtime test that returns a value. These metrics can be associated with physical resources such as servers

and switches, virtual resources such as VMs and containers, or high-level abstractions such as the *Connectivity Service* described in the next section. Given these many possible sources of data, the job of the metrics monitoring stack is to collect, archive, visualize, and optionally analyze this data. Prometheus is a popular open source tool for storing and querying metrics.

Further Reading

[Prometheus](#).

Logs are the qualitative data that is generated whenever a noteworthy event occurs. This information can be used to identify problematic operating conditions (i.e., it may trigger an alert), but more commonly, it is used to troubleshoot problems after they have been detected. Various system components—all the way from the low-level OS kernel to high-level cloud services—write messages that adhere to a well-defined format to the log. These messages include a timestamp, which makes it possible for the logging stack to parse and correlate messages from different components. Elasticsearch is a widely-used tool for storing and analyzing log messages.

Further Reading

[ElasticSearch](#).

Traces are a record of causal relationships (e.g., Service A calls Service B) resulting from user-initiated transactions or jobs. They are related to logs, but provide more specialized information about the context in which different events happen. Tracing is well understood in a single program, where an execution trace is commonly recorded as an in-memory call stack, but traces are inherently distributed across a graph of network-connected microservices in a cloud setting. This makes the problem challenging, but also critically important because it is often the case that the only way to understand time-dependent phenomena—such as why a particular resource is overloaded—is to understand how multiple independent workflows interact with each other. Jaeger is a popular open source tool used for tracing.

Further Reading

[Jaeger: End-to-End Distributed Tracing](#).

Finally, note that our framing of monitoring and telemetry as part of the online aspect of management is somewhat simplistic. Certainly telemetry data is collected from online processes embedded in a running system, and such data can be coupled with online control operations to realize closed-loop control, but it is also the case that some telemetry data is evaluated offline. This is true for logs and traces used to diagnose problems, and for performance data used to make provisioning decisions, both of which can lead to code changes that feed back into the next iteration of lifecycle management.

6.4 Connectivity API

The visible aspect of a 5G service is the programmatic interface it provides to users, giving them the ability to control and customize the underlying connectivity service. This API is implemented by the Service Orchestrator outlined in the previous section, but what we really care about is the interface itself. Using Aether as a concrete example, this section describes such an API.

Like many cloud services, the API for 5G-as-a-Service is RESTful. This means it supports REST's GET, POST, PATCH, and DELETE operations on a set of resources (objects):

- GET: Retrieve an object.
- POST: Create an object.
- PUT, PATCH: Modify an existing object.
- DELETE: Delete an object.

Each object, in turn, is typically defined by a data model. In Aether this model is specified in YANG, but rather than dive into the particulars of YANG, this section describes the models informally by describing the relevant fields.

Every object contains an *id* field that is used to uniquely identify the object. Some objects contain references to other objects. For

example, many objects contain references to the *Enterprise* object, which allows them to be associated with a particular enterprise. That is, references are constructed using the *id* field of the referenced object.

In addition to the *id* field, several other fields are also common to all models. These include:

- *description*: A human-readable description, used to store additional context about the object.
- *display-name*: A human-readable name that is shown in the GUI.

As these fields are common to all models, we omit them from the per-model descriptions that follow. Note that we use upper case to denote a model (e.g., *Enterprise*) and lower case to denote a field within a model (e.g., *enterprise*).

6.4.1 *Enterprises*

Aether is deployed in enterprises, and so needs to define a representative set of organizational abstractions. These include *Enterprise*, which forms the root of a customer-specific hierarchy. The *Enterprise* model is referenced by many other objects, and allows those objects to be scoped to a particular Enterprise for ownership and role-based access control purposes. *Enterprise* contains the following fields:

- *connectivity-service*: A list of backend subsystems that implement connectivity for this enterprise. This list corresponds to the API endpoint for the SD-Core, SD-Fabric, and SD-RAN components.

Enterprises are further divided into *Sites*. A site is a point-of-presence for an *Enterprise* and may be either physical or logical (i.e., a single geographic location could contain several logical sites). The *Site* model, in turn, contains the following fields:

- *enterprise*: A link to the *Enterprise* that owns this site.

- *imsi-definition*: A description of how IMSIs are constructed for this site. It consists of the following sub-fields:
 - *mcc*: Mobile country code.
 - *mnc*: Mobile network code.
 - *enterprise*: A numeric enterprise id.
 - *format*: A mask that defines how the above three fields are encoded in an IMSI. For example *CCCNNEEESSSSSS* specifies an IMSI with a 3-digit MCC, a 3-digit MNC, a 3-digit ENT, and a 6-digit subscriber.

As a reminder, an IMSI is burned into every SIM card, and is used to identify and locate UEs throughout the global cellular network.

6.4.2 Slices

Aether models 5G connectivity as a *Slice*, which represents an isolated communication channel (and associated QoS parameters) that connects a set of devices (modeled as a *Device-Group*) to a set of applications (each of which is modeled as an *Application*). For example, an enterprise might configure one slice to carry IoT traffic and another slice to carry video traffic. The *Slice* model has the following fields:

- *device-group*: A list of *Device-Group* objects that can participate in this *Slice*. Each entry in the list contains both the reference to the *Device-Group* as well as an *enable* field which may be used to temporarily remove access to the group.
- *application*: A list of *Application* objects that are either allowed or denied for this *Slice*. Each entry in the list contains both a reference to the *Application* as well as an *allow* field which can be set to *true* to allow the application or *false* to deny it.
- *template*: Reference to the *Template* that was used to initialize this *Slice*.
- *upf*: Reference to the User Plane Function (*UPF*) that should be used to process packets for this *Slice*. Multiple *Slices* may share a single *UPF*.

- *enterprise*: Reference to the *Enterprise* that owns this *Slice*.
- *site*: Reference to the *Site* where this *Slice* is deployed.
- *sst, sd*: 3GPP-defined slice identifiers assigned by the operations team.
- *mbr.uplink, mbr.downlink, mbr.uplink-burst-size, mbr.downlink-burst-size*: Maximum bit-rate and burst sizes for this slice.

The rate-related parameters are initialized using a selected *template*, as described below, but these values may be changed at runtime. Also note that this example illustrates how modeling can be used to enforce invariants, in this case, that the *Site* of the *UPF* and *Device-Group* must match the *Site* of the *Slice*. That is, the physical devices that connect to a slice and the *UPF* that implements the core segment of the slice must be constrained to a single physical location.

At one end of a *Slice* is a *Device-Group*, which identifies a set of devices that are allowed to use the *Slice* to connect to various applications. The *Device-Group* model contains the following fields:

- *imsis*: A list of IMSI ranges. Each range has the following fields:
 - *name*: Name of the range. Used as a key.
 - *imsi-range-from*: First IMSI in the range.
 - *imsi-range-to*: Last IMSI in the range.
Can be omitted if the range only contains one IMSI.
- *ip-domain*: Reference to an *IP-Domain* object that describes the IP and DNS settings for UEs within this group.
- *site*: Reference to the site where this *Device-Group* may be used. (This field indirectly identifies the *Enterprise* since a *Site* contains a reference to *Enterprise*.)
- *mbr.uplink, mbr.downlink*: Maximum bit-rate for the device group.
- *traffic-class*: The traffic class to be used for devices in this

group.

At the other end of a *Slice* is a list of *Application* objects, which specifies the endpoints for the program devices talk to. The *Application* model contains the following fields:

- *address*: The DNS name or IP address of the endpoint.
- *endpoint*: A list of endpoints. Each has the following fields:
 - *name*: Name of the endpoint. Used as a key.
 - *port-start*: Starting port number.
 - *port-end*: Ending port number.
 - *protocol*: Protocol (*TCP|UDP*) for the endpoint.
 - *mbr.uplink*, *mbr.downlink*: Maximum bitrate for devices communicating with this application.
 - *traffic-class*: Traffic class for devices communicating with this application.
- *enterprise*: Link to an *Enterprise* object that owns this application. May be left empty to indicate a global application that may be used by multiple enterprises.

Note that Aether’s *Slice* abstraction is similar to 3GPP’s specification of a “slice”, but the *Slice* model includes a combination of 3GPP-specified identifiers (e.g., *sst* and *sd*), and details about the underlying implementation (e.g., *upf* denotes the UPF implementation for the Core’s user plane). The *Slice* model also includes fields related to RAN slicing, with the Service Orchestrator responsible for stitching together end-to-end connectivity across the RAN, Core, and Fabric.

6.4.3 QoS Profiles

Associated with each *Slice* is a QoS-related profile that governs how traffic carried by that slice is to be treated. This starts with a

Template model, which defines the valid (accepted) connectivity settings. The Aether Operations team is responsible for defining these (the features they offer must be supported by the backend subsystems), with enterprises selecting the template they want applied to any instances of the connectivity service they create (e.g., via a drop-down menu). That is, templates are used to initialize *Slice* objects. The *Template* model has the following fields:

- *sst*, *sd*: Slice identifiers, as specified by 3GPP.
- *mbr.uplink*, *mbr.downlink*: Maximum uplink and downlink bandwidth.
- *mbr.uplink-burst-size*, *mbr.downlink-burst-size*: Maximum burst size.
- *traffic-class*: Link to a *Traffic-Class* object that describes the type of traffic.

You will see that the *Device-Group* and *Application* models include similar fields. The idea is that QoS parameters are established for the slice as a whole (based on the selected *Template*) and then individual devices and applications connected to that slice can define their own, more-restrictive QoS parameters on an instance-by-instance basis.

Finally, the *Traffic-Class* model specifies the classes of traffic, and includes the following fields:

- *arp*: Allocation and retention priority.
- *qci*: QoS class identifier.
- *pelr*: Packet error loss rate.
- *pdb*: Packet delay budget.

6.4.4 Other Models

The above description references other models, which we do not fully describe here. They include *AP-List*, which specifies a list of access points (radios); *IP-Domain*, which specifies IP and DNS settings; and *UPF*, which specifies the User Plane Function (the data plane element of the SD-Core) that is to forward packets on behalf of this particular instance of the connectivity service. The *UPF*

model is necessary because Aether supports two different implementations: one runs as a microservice on a server and the other runs as a P4 program loaded into the switching fabric. Both implementations are described in Chapter 5.



About The Book

Source for *Private 5G: A Systems Approach* is available on GitHub under terms of the [Creative Commons \(CC BY-NC-ND 4.0\)](#) license. The community is invited to contribute corrections, improvements, updates, and new material under the same terms. While this license does not automatically grant the right to make derivative works, we are keen to discuss derivative works (such as translations) with interested parties. Please reach out to discuss@systemsapproach.org.

If you make use of this work, the attribution should include the following information:

Title: Private 5G: A Systems Approach

Authors: Larry Peterson, Oguz Sunay, and Bruce Davie

Source: <https://github.com/SystemsApproach/private5g>

License: CC BY-NC-ND 4.0

This book incorporates introductory and background content from:

Title: 5G Mobile Networks: A Systems Approach

Authors: Larry Peterson and Oguz Sunay

Source: <https://github.com/SystemsApproach/5g>

License: CC BY-NC-ND 4.0

a version of which was also published by Morgan & Claypool in 2020 as part of their Synthesis Lectures on Network Systems.

Designations used by companies to distinguish their products are often claimed as trademarks or registered trademarks. In all instances in which Systems Approach, LLC, is aware of a claim, the

product names appear in initial capital or all capital letters. Readers, however, should contact the appropriate companies for more complete information regarding trademarks and registration.

Read the Book

This book is part of the [Systems Approach Series](#), with an online version published at <https://5G.systemsapproach.org>.

For those users looking for our earlier 5G book, *5G Mobile Networks: A Systems Approach*, published as this URL, you can still find source for it archived on [GitHub](#). This new book incorporates the background material covered in the old one, plus goes into significant detail about how Private 5G is implemented and deployed as a managed cloud service.

To track progress and receive notices about new versions, you can follow the project on [Facebook](#) and [Mastodon](#). To read a running commentary on how the Internet is evolving, follow the [Systems Approach on Substack](#).

Build the Book

To build a web-viewable version, you first need to download the source:

```
$ mkdir ~/systemsapproach
$ cd ~/systemsapproach
$ git clone https://github.com/SystemsApproach/private5g
$ cd private5g
```

The build process is stored in the Makefile and requires Python be installed. The Makefile will create a virtualenv (`venv-docs`) which installs the documentation generation toolset. You may also need to install the `enchant` C library using your system's package manager for the spelling checker to function properly.

To generate HTML in `_build/html`, run `make html`.

To check the formatting of the book, run `make lint`.

To check spelling, run `make spelling`. If there are additional words, names, or acronyms that are correctly spelled but not in the dictionary, please add them to the `dict.txt` file.

To see the other available output formats, run `make`.

Contribute to the Book

We hope that if you use this material, you are also willing to contribute back to it. If you are new to open source, you might check out this [How to Contribute to Open Source](#) guide. Among other things, you'll learn about posting *Issues* that you'd like to see addressed, and issuing *Pull Requests* to merge your improvements back into GitHub.

If you'd like to contribute and are looking for something that needs attention, see the [wiki](#) for the current TODO list.

About The Authors

Larry Peterson is the Robert E. Kahn Professor of Computer Science, Emeritus at Princeton University, where he served as Chair from 2003-2009. His research focuses on the design, implementation, and operation of Internet-scale distributed systems, including the widely used PlanetLab and MeasurementLab platforms. He is currently contributing to the Aether access-edge cloud project at the Open Networking Foundation (ONF), where he serves as Chief Scientist. Peterson is a member of the National Academy of Engineering, a Fellow of the ACM and the IEEE, the 2010 recipient of the IEEE Kobayashi Computer and Communication Award, and the 2013 recipient of the ACM SIGCOMM Award. He received his Ph.D. degree from Purdue University.

Oguz Sunay is Chief Architect for the Network and Edge Group (NEX) at Intel, which he joined as part of Intel's acquisition of the Open Networking Foundation (ONF) engineering team. While at ONF, he was Vice President for Research & Development, where led all mobile-related projects. Before joining ONF, Sunay was the CTO at Argela-USA, where he was the innovator of a Programmable Radio Access Network Architecture (ProGRAN) for 5G that enabled the world's first dynamically programmable RAN slicing solution. He has also held prior industry positions at Nokia Research Center and Bell Laboratories, where he focused on 3G and 4G end-to-end systems architectures and participated and chaired various standardization activities. Sunay has also spent over 10 years in academia, as a Professor of Electrical and Computer Engineering. He holds many U.S. and European patents on various aspects of 3G, 4G, and 5G, and has authored numerous journal and conference publications. He received his Ph.D. and M.Sc. from Queen's University, Canada, and his B.Sc.Hon. from METU, Turkey.

Bruce Davie is a computer scientist noted for his contributions to the field of networking. He is a former VP and CTO for the Asia Pacific region at VMware. He joined VMware during the acquisition of Software Defined Networking (SDN) startup Nicira. Prior to that, he was a Fellow at Cisco Systems, leading a team of architects responsible for Multiprotocol Label Switching (MPLS). Davie has over 30 years of networking industry experience and has co-authored 17 RFCs. He was recognized as an ACM Fellow in 2009 and chaired ACM SIGCOMM from 2009 to 2013. He was also a visiting lecturer at the Massachusetts Institute of Technology for five years. Davie is the author of multiple books and the holder of more than 40 U.S. Patents.

Read the Latest!

Substack Newsletter: Stay up-to-date with the latest developments by following the authors on [Substack](#), where they connect the concepts and lessons in this book to what's happening in the Internet today.

Book Series: Also check out our companion books that cover emerging topics in more depth.

- [TCP Congestion Control: A Systems Approach](#)
- [Edge Cloud Operations: A Systems Approach](#)
- [Software-Defined Networks: A Systems Approach](#)

Print Copies

We make all books in the *Systems Approach* series available as both print and e-books. This book is available via Amazon: [Private 5G: A Systems Approach](#)

Book Series: Also check out our companion books that cover networking and emerging topics in more depth.

- [Computer Networks: A Systems Approach](#)
- [TCP Congestion Control: A Systems Approach](#)
- [Edge Cloud Operations: A Systems Approach](#)
- [Software-Defined Networks: A Systems Approach](#)

As participants in the Amazon Associate program we may earn income from qualifying purchases using the links above.

Overview

Many of the implementation details presented in this book were informed by Aether, an open source 5G edge cloud connectivity service. A multi-site deployment of Aether has been running since 2020 in support of the *Pronto Project*, but that deployment depends on an ops team with significant insider knowledge about Aether's engineering details. It is difficult for others to reproduce that know-how and bring up their own Aether clusters.

[Aether OnRamp](#) is a re-packaging of Aether to address that problem. It provides an incremental path for users to:

- Learn about and observe all the moving parts in Aether.
- Customize Aether for different target environments.
- Experiment with scalable edge communication.
- Deploy and operate Aether with live traffic.

Aether OnRamp begins with a *Quick Start* deployment that is easy to bring up in a single VM, but then goes on to prescribe a sequence of steps users can follow to deploy increasingly complex configurations. OnRamp refers to each such configuration as a *blueprint*, and the set supports both emulated and physical RANs, along with the runtime machinery needed to operate an Aether cluster supporting live 5G workloads. (OnRamp also defines a 4G blueprint that can be used to connect one or more physical eNBs, but we postpone a discussion of that capability until a later section. Everything else in this guide assumes 5G.)

Note that Aether OnRamp does not include SD-Fabric, which depends on programmable switching hardware. Readers interested in learning more about that capability (including a P4-based UPF) should see the Hands-on Programming appendix of our companion SDN book.

Further Reading

[Pronto Project: Building Secure Networks Through Verifiable Closed-Loop Control](#).

[Hands-on Programming \(Appendix\). Software-Defined Networks: A Systems Approach](#). November 2021.

Aether is a complex system, assembled from multiple components spanning several Git repositories. These include repos for different subsystems (e.g., AMP, SD-Core, SD-RAN), but also for different stages of the development pipeline (e.g., source code, deployment artifacts, configuration specs). This rest of this section identifies all the Aether-related repositories, with the OnRamp repos listed at the end serving as the starting point for anyone who wants to come up to speed on the rest of the system.

Troubleshooting Hint

This guide includes *Troubleshooting Hints* like this one. Our first hint is to recommend that the guide be followed sequentially. This is because each section establishes a milestone that may prove useful when you find yourself trying to troubleshoot a problem in a later section. For example, isolating a problem with a physical gNB is easier if you know that connectivity to the AMF and UPF works correctly, which the *Emulated RAN* section helps to establish.

Our second hint is to join the `#aether-onramp` channel of the [ONF Workspace](#) on Slack, where questions about using OnRamp to bring up Aether are asked and answered. The `Troubleshooting` bookmark for that channel includes summaries of known issues.

Source Repos

Source code for Aether and all of its subsystems can be found in the following repositories:

- Gerrit repository for the CORD Project (<https://gerrit.opencord.org>): Microservices for AMP, plus source for the jobs that implement the CI/CD pipeline.

- GitHub repository for the OMEC Project (<https://github.com/omec-project>): Microservices for SD-Core, plus the emulator (gNBsim) that subjects SD-Core to RAN workloads.
- GitHub repository for the ONOS Project (<https://github.com/onosproject>): Microservices for SD-Fabric and SD-RAN, plus the YANG models used to generate the Aether API.
- GitHub repository for the Stratum Project (<https://github.com/stratum>): On-switch components of SD-Fabric.

For Gerrit, you can either browse Gerrit (select the *master* branch) or clone the corresponding `<repo-name>` by typing:

```
$ git clone ssh://gerrit.opencord.org:29418/<repo-name>
```

If port 29418 is blocked by your network administrator, you can try cloning using https instead of ssh:

```
$ git clone https://gerrit.opencord.org/<repo-name>
```

Anyone wanting to participate in Aether's ongoing development will want to learn how to contribute new features to these source repos.

Artifact Repos

Aether includes a *Continuous Integration (CI)* pipeline that builds deployment artifacts (e.g., Helm Charts, Docker Images) from the source code. These artifacts are stored in the following repositories:

Helm Charts

<https://charts.aetherproject.org>
<https://charts.onosproject.org>
<https://charts.opencord.org>
<https://charts.atomix.io>
<https://sdrancharts.onosproject.org>
<https://charts.rancher.io/>

Docker Images

<https://registry.aetherproject.org>

Note that as of version 1.20.8, Kubernetes uses the [Containerd](#) runtime system instead of Docker. This is transparent to anyone using Aether, which manages containers indirectly through Kubernetes (e.g., using `kubectl`), but does impact anyone who directly depends on the Docker toolchain. Also note that while Aether documentation often refers its use of “Docker containers,” it is now more accurate to say that Aether uses [OCI-Compliant containers](#).

The Aether CI pipeline keeps the above artifact repos in sync with the source repos listed above. Among those source repos are the source files for all the Helm Charts:

ROC: <https://gerrit.opencord.org/plugins/gitiles/roc-helm-charts>

SD-RAN: <https://github.com/onosproject/sdran-helm-charts>

SD-Core: <https://gerrit.opencord.org/plugins/gitiles/sdcore-helm-charts>

SD-Fabric (Servers): <https://github.com/onosproject/onos-helm-charts>

SD-Fabric (Switches): <https://github.com/stratum/stratum-helm-charts>

The QA tests run against code checked into these source repos can be found here:

<https://gerrit.opencord.org/plugins/gitiles/aether-system-tests>

The specification for the CI pipeline, which invokes these QA tests, gates merge requests, and publishes artifacts, can be found here:

<https://gerrit.opencord.org/plugins/gitiles/aether-ci-management>

For more information about Aether’s CI pipeline, including its QA and version control strategy, we recommend the Lifecycle Management chapter of our companion Edge Cloud Operations

book.

Further Reading

L. Peterson, A. Bavier, S. Baker, Z. Williams, and B. Davie. [Edge Cloud Operations: A Systems Approach](#). June 2022.

OnRamp Repos

The process to deploy the artifacts listed above, sometimes referred to as GitOps, manages the *Continuous Deployment (CD)* half of the CI/CD pipeline. OnRamp's approach to GitOps uses a different mechanism than the one the ONF ops team originally used to manage its multi-site deployment of Aether. The latter approach has a large startup cost, which has proven difficult to replicate. (It also locks you into deployment toolchain that may or may not be appropriate for your situation.)

In its place, OnRamp adopts minimal Ansible tooling. This makes it easier to take ownership of the configuration parameters that define your specific deployment scenario. The rest of this guide walks you through a step-by-step process of deploying and operating Aether on your own hardware. For now, we simply point you at the collection of OnRamp repos:

Deploy Aether: <https://github.com/opennetworkinglab/aether-onramp>
Deploy 5G Core: <https://github.com/opennetworkinglab/aether-5gc>
Deploy 4G Core: <https://github.com/opennetworkinglab/aether-4gc>
Deploy Management Plane: <https://github.com/opennetworkinglab/aether-amp>
Deploy 5G RAN Simulator: <https://github.com/opennetworkinglab/aether-gnbsim>
Deploy Kubernetes: <https://github.com/opennetworkinglab/aether-k8s>

It is the first repo that defines a way to integrate all of the Aether artifacts into an operational system. That repo, in turn, includes the

other repos as submodules. Note that each of the submodules is self-contained if you are interested in deploying just that subsystem, but this guide approaches the deployment challenge from an integrated, end-to-end perspective.

Because OnRamp uses Ansible as its primary deployment tool, a general understanding of Ansible is helpful (see the suggested reference). However, this guide walks you, step-by-step, through the process of deploying and operating Aether, so previous experience with Ansible is not a requirement. Note that Ansible has evolved to be both a “Community Toolset” anyone can use to manage a software deployment, and an “Automation Platform” offered as a service by RedHat. OnRamp uses the toolset, but not the platform/service.

Further Reading

[Overview: How Ansible Works.](#)

Aether OnRamp is still a work in progress, but anyone interested in participating in that effort is encouraged to join the discussion on Slack in the [ONF Community Workspace](#). A roadmap for the work that needs to be done can be found in the [Aether OnRamp Wiki](#).

Add ▾ Remove Edit Disk Action ▾ Revert		
	Memory	70.00 GiB
	Processors	32 (1 sockets, 32 cores) [host]
	BIOS	Default (SeaBIOS)
	Display	Default
	Machine	Default (i440fx)
	SCSI Controller	VirtIO SCSI single
	CD/DVD Drive (ide2)	local:iso/ubuntu-20.04.5-live-server-amd64.iso,media=cdrom,size=1373568K
	Hard Disk (scsi0)	local-lvm:vm-160-disk-0,iothread=1,size=64G
	Network Device (net0)	virtio=9E:37:E6:F6:4A:99,bridge=evpn1gnt,firewall=1

Quick Start

This section describes a low-overhead way to get started with OnRamp. It brings up a one-node Kubernetes cluster, deploys a 5G version of SD-Core on that cluster, and runs an emulated 5G workload against the 5G Core. It assumes a low-end server that meets the following requirements:

- Haswell CPU (or newer), with at least 4 CPU cores and 16GB RAM.
- Clean install of Ubuntu 20.04 or 22.04, with 5.15 (or later) kernel.

For example, something like an Intel NUC is likely more than enough to get started.

While it's possible to use OnRamp to deploy Aether on a laptop (e.g., in a VirtualBox VM), because the goal is to eventually scale a deployment and/or run Aether 24/7, OnRamp has been developed and tested on physical servers and server-based VMs. The latter includes Proxmox (see the example configuration shown in [Figure 47](#)); AWS (specify a `t2.xlarge` instance); and CloudLab (specify either a `small-lan` or `single-pc-ubuntu` instance).

Figure 47. Example configuration of Proxmox VM.

Prep Environment

To install Aether OnRamp, you must be able to run `sudo` without a password, and there should be no firewall running on the server. You can verify this is the case by executing the following, which should report `Status: inactive`:

```
$ sudo ufw status
Status: inactive
```

Your server should use *systemd-networkd* to configure the network, which you can verify by typing:

```
$ systemctl status systemd-networkd.service
```

Note that Aether assumes Ubuntu Server (as opposed to Ubuntu Desktop), the main implication being that it uses *systemd-networkd* rather than *Network Manager* to manage network settings. It is possible to work around this requirement, but be aware that doing so may impact the Ansible playbook for installing SD-Core.

OnRamp depends on Ansible, which you can install on your server as follows:

```
$ sudo apt install pipx
$ sudo apt install python3.8-venv
$ pipx install --include-deps ansible
$ pipx ensurepath
$ sudo apt-get install sshpass
```

Once installed, displaying the Ansible version number should result in output similar to the following:

```
$ ansible --version
ansible [core 2.11.12]
  config file = None
  configured module search path = ['/home/foo/.ansible/p
  ansible python module location = /home/foo/.local/lib/
  ansible collection location = /home/foo/.ansible/colle
  executable location = /home/foo/.local/bin/ansible
```

```
python version = 3.6.9 (default, Mar 10 2023, 16:46:00)
jinja version = 3.0.3
libyaml = True
```

Note that a fresh install of Ubuntu may be missing other packages that you need (e.g., `git`, `curl`, `make`), but you will be prompted to install them as you step through the Quick Start sequence.

Download Aether OnRamp

Once ready, clone the Aether OnRamp repo on this target deployment server:

```
$ git clone --recursive https://github.com/opennetworking/aether-onramp
$ cd aether-onramp
```

Taking a quick look at your `aether-onramp` directory, there are four things to note:

1. The `deps` directory contains the Ansible deployment specifications for all the Aether subsystems. Each of these subdirectories (e.g., `deps/5gc`) is self-contained, meaning you can execute the Make targets in each individual directory. Doing so causes Ansible to run the corresponding playbook. For example, the installation playbook for the 5G Core can be found in `deps/5gc/roles/core/tasks/install.yml`.
2. The Makefile in the main OnRamp directory imports (`#include`) the per-subsystem Makefiles, meaning all the individual steps required to install Aether can be managed from this main directory. The Makefile includes comments listing the key Make targets defined by the included Makefiles. *Importantly, the rest of this guide assumes you are working in the main OnRamp directory, and not in the individual subsystems.*
3. File `vars/main.yml` defines all the Ansible variables you will potentially need to modify to specify your deployment scenario. This file is the union of all the per-component `var/main.yml` files you find in the corresponding `deps`

directory. This top-level variable file overrides the per-component var files, so you will not need to modify the latter. Note that the `vars` directory contains several variants of `main.yml`, where we think of each as specifying a *blueprint* for a different configuration of Aether. The default `main.yml` (which is the same as `main-quickstart.yml`) gives the blueprint for the Quick Start deployment described in this section; we'll substitute the other blueprints in later sections.

4. File `hosts.ini` (host inventory) is Ansible's way of specifying the set of servers (physical or virtual) that Ansible targets with various installation playbooks. The default version of `hosts.ini` included with OnRamp is simplified to run everything on a single server (the one you've cloned the repo onto), with additional lines you may eventually need for a multi-node cluster commented out.

Set Target Parameters

The Quick Start deployment described in this section requires that you modify two sets of parameters to reflect the specifics of your target deployment.

The first set is in file `hosts.ini`, where you will need to give the IP address and login credentials for the server you are working on. At this stage, we assume the server you downloaded OnRamp onto is the same server you will be installing Aether on.

```
node1  ansible_host=10.76.28.113 ansible_user=aether ans
```

In this example, address `10.76.28.113` and the three occurrences of the string `aether` need to be replaced with the appropriate values. Note that if you set up your server to use SSH keys instead of passwords, then `ansible_password=aether` needs to be replaced with `ansible_ssh_private_key_file=~/.ssh/id_rsa` (or wherever your private key can be found).

The second set of parameters is in `vars/main.yml`, where the **two** lines currently reading

```
data_iface: ens18
```

need to be edited to replace `ens18` with the device interface for your server, and the line specifying the IP address of the Core's AMF needs to be edited to reflect your server's IP address:

```
amf:  
  ip: "10.76.28.113"
```

You can learn your server's IP address and interface using the Linux `ip` command:

```
$ ip a  
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN  
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00  
    inet 127.0.0.1/8 scope host lo  
        valid_lft forever preferred_lft forever  
    inet6 ::1/128 scope host  
        valid_lft forever preferred_lft forever  
2: ens18: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fqdisc  
    link/ether 2c:f0:5d:f2:d8:21 brd ff:ff:ff:ff:ff:ff  
    inet 10.76.28.113/24 metric 100 brd 10.76.28.255 scope global  
        valid_lft forever preferred_lft forever  
    inet6 fe80::2ef0:5dff:fef2:d821/64 scope link  
        valid_lft forever preferred_lft forever
```

In this example, the reported interface is `ens18` and the IP address is `10.76.28.113` on subnet `10.76.28.0/24`. We will use these three values as a running example throughout the guide, as a placeholder for your local details.

Troubleshooting Hint

Due to a limitation in gNBsim (the RAN emulator introduced later in this section), it is necessary for your server to be configured with IPv6 enabled (as the `inet6` line in the example output indicates is the case for interface `ens18`). If IPv6 is not enabled, the emulated RAN will not successfully connect to the AMF.

Note that `vars/main.yml` and `hosts.ini` are the only two files you need to modify for now, but there are additional config files

that you may want to modify as we move beyond the Quick Start deployment. We'll identify those files throughout this section, for informational purposes, and revisit them in later sections.

Many of the tasks specified in the various Ansible playbooks result in calls to Kubernetes, either directly via `kubectl`, or indirectly via `helm`. This means that after executing the sequence of Makefile targets described in the rest of this guide, you'll want to run some combination of the following commands to verify that the right things happened:

```
$ kubectl get pods --all-namespaces
$ helm repo list
$ helm list --namespace kube-system
```

The first reports the set of Kubernetes namespaces currently running; the second shows the known set of repos you are pulling charts from; and the third shows the version numbers of the charts currently deployed in the `kube-system` namespace.

If you are not familiar with `kubectl` (the CLI for Kubernetes), we recommend that you start with [Kubernetes Tutorial](#).

Note that we have not yet installed Kubernetes or Helm, so these commands are not yet available. At this point, the only verification step you can take is to type the following:

```
$ make aether-pingall
```

The output should show that Ansible is able to securely connect to all the nodes in your deployment, which is currently just the one that Ansible knows as `node1`.

Install Kubernetes

The next step is to bring up an RKE2.0 Kubernetes cluster on your target server. Do this by typing:

```
$ make aether-k8s-install
```

Note that the Ansible playbooks triggered by this (and other) make targets will output red results from time-to-time (indicating an exception or failure), but as long as Ansible keeps progressing through the playbook, such output can be safely ignored.

Once the playbook completes, executing `kubectl` will show the `kube-system` namespace running, with output looking something like the following:

```
$ kubectl get pods --all-namespaces
NAMESPACE      NAME
kube-system    cloud-controller-manager-node1
kube-system    etcd-node1
kube-system    helm-install-rke2-canal-8s67r
kube-system    helm-install-rke2-coredns-bk5rh
kube-system    helm-install-rke2-ingress-nginx-lsjz2
kube-system    helm-install-rke2-metrics-server-t8kxf
kube-system    helm-install-rke2-multus-tbbhc
kube-system    kube-apiserver-node1
kube-system    kube-controller-manager-node1
kube-system    kube-multus-ds-96cnl
kube-system    kube-proxy-node1
kube-system    kube-scheduler-node1
kube-system    rke2-canal-h79qq
kube-system    rke2-coredns-rke2-coredns-869b5d56d4-tffjh
kube-system    rke2-coredns-rke2-coredns-autoscaler-5b947
kube-system    rke2-ingress-nginx-controller-s68rx
kube-system    rke2-metrics-server-6564db4569-snnv4
```

If you are interested in seeing the details about how Kubernetes is customized for Aether, look at `deps/k8s/roles/rke2/templates/master-config.yaml`. Of particular note, we have instructed Kubernetes to allow service for ports ranging from 2000 to 36767 and we are using the `multus` and `canal` CNI plugins.

Install SD-Core

We are now ready to bring up the 5G version of the SD-Core. To do

that, type:

```
$ make aether-5gc-install
```

kubectl will now show the omec namespace running (in addition to kube-system), with output similar to the following:

```
$ kubectl get pods -n omec
```

NAME	READY	STATUS
amf-5887bbf6c5-pc9g2	1/1	Running
ausf-6dbb7655c7-42z7m	1/1	Running
kafka-0	1/1	Running
metricfunc-b9f8c667b-r2x9g	1/1	Running
mongodb-0	1/1	Running
mongodb-1	1/1	Running
mongodb-arbiter-0	1/1	Running
nrf-54bf88c78c-kcm7t	1/1	Running
nssf-5b85b8978d-d29jm	1/1	Running
pcf-758d7cfb48-dwz9x	1/1	Running
sd-core-zookeeper-0	1/1	Running
simapp-6cccd6f787-jnxc7	1/1	Running
smf-7f89c6d849-wzqvx	1/1	Running
udm-768b9987b4-9qz4p	1/1	Running
udr-8566897d45-kv6zd	1/1	Running
upf-0	5/5	Running
webui-5894ffd49d-gg2jh	1/1	Running

If you see problematic pods that are not getting into the Running state, a reset usually corrects the problem. Type:

```
make aether-resetcore
```

Once running, you will recognize pods that correspond to many of the microservices discussed in [Chapter 5](#). For example, amf-5887bbf6c5-pc9g2 implements the AMF. Note that for historical reasons, the Aether Core is called omec instead of sd-core.

Troubleshooting Hint

If you see failures of the `find ens18's netplan network directory` task in the `router` role, it indicates that *systemd-networkd* is not configured as expected. Check the [OnRamp Troubleshooting Wiki Page](#) for possible workarounds.

If you are interested in seeing the details about how SD-Core is configured, look at `deps/5gc/roles/core/templates/radio-5g-values.yaml`. This is an example of a *values override* file that Helm passes along to Kubernetes when launching the service. Most of the default settings will remain unchanged, with the main exception being the `subscribers` block of the `omec-sub-provision` section. This block will eventually need to be edited to reflect the SIM cards you actually deploy. We return to this topic in the section describing how to bring up a physical gNB.

Run Emulated RAN Test

We can now test SD-Core with emulated traffic by typing:

```
$ make aether-gnbsim-install
$ make aether-gnbsim-run
```

Note that you can re-execute the `aether-gnbsim-run` target multiple times, where the results of each run are saved in a file within the Docker container running the test. You can access that file by typing:

```
$ docker exec -it gnbsim-1 cat summary.log
```

If successful, the output should look like the following:

```
2023-08-09T19:57:09Z [INFO] [GNBSIM] [Summary] Profile Name
2023-08-09T19:57:09Z [INFO] [GNBSIM] [Summary] UEs Passed:
2023-08-09T19:57:09Z [INFO] [GNBSIM] [Summary] Profile Sta
```

This particular test, which runs the cryptically named `pdusessest` profile, emulates five UEs, each of which: (1) registers with the Core, (2) initiates a user plane session, and (3) sends a minimal data packet over that session. In addition to displaying the summary results, you can also open a shell in the `gnbsim-1`

container, where you can view the full trace of every run of the emulation, each of which has been saved in a timestamped file:

```
$ docker exec -it gnbsim-1 bash
bash-5.1# ls
gnbsim                                gnbsim1-20230809T125702.
gnbsim.log                            gnbsim1-20230809T125702.
bash-5.1# more gnbsim1-20230809T125702.log
2023-08-09T19:57:05Z [INFO] [GNBSIM] [App] App Name: GNBSIM
2023-08-09T19:57:05Z [INFO] [GNBSIM] [App] Setting log lev
2023-08-09T19:57:05Z [INFO] [GNBSIM] [GNodeB] [gnb1] GNodeB
2023-08-09T19:57:05Z [INFO] [GNBSIM] [GNodeB] [UserPlaneTra
2023-08-09T19:57:05Z [INFO] [GNBSIM] [GNodeB] Current rang
2023-08-09T19:57:05Z [INFO] [GNBSIM] [GNodeB] Current ID r
2023-08-09T19:57:05Z [INFO] [GNBSIM] [GNodeB] [ControlPlane
...
```

Troubleshooting Hint

If `summary.log` is empty, it means the emulation did not run due to a configuration error. To debug the problem, open a bash shell on the `gnbsim` container (as shown in the preceding example), and look at `gnbsim.log`. Output that includes `failed to connect amf` and `err: address family not supported by protocol` indicates that your server does not have IPv6 enabled.

Troubleshooting Hint

If `summary.log` reports `UEs Passed: 0` , `UEs Failed: 5` then it may be the case that SD-Core did not come up cleanly. Type `make aether-resetcore`, and after verifying all pods are running with `kubectl`, run `gnbsim` again.

Another possibility is that you have multiple SD-Cores running in the same broadcast domain. This causes ARP to behave in unexpected ways, which interferes with OnRamp's ability to establish a route to the UPF pod.

If you are interested in the config file that controls the test, including the option of enabling other profiles, take a look at `deps/gnbsim/config/gnbsim-default.yaml`. We return to

the issue of customizing gNBsim in a later section, but for now there are some simple modifications you can try. For example, the following code block defines a set of parameters for `pdusesest` (also known as `profile2`):

```
- profileType: pdusesest           # UE Initiated Session
profileName: profile2
enable: true
gnbName: gnb1
execInParallel: false
startImsi: 208930100007487
ueCount: 5
defaultAs: "192.168.250.1"
perUserTimeout: 100
plmnId:
  mcc: 208
  mnc: 93
dataPktCount: 5
opc: "981d464c7c52eb6e5036234984ad0bcf"
key: "5122250214c33e723a5dd523fc145fc0"
sequenceNumber: "16f3b3f70fc2"
```

You can edit `ueCount` to change the number of UEs included in the emulation (currently limited to 100) and you can set `execInParallel` to `true` to emulate those UEs connecting to the Core in parallel (rather than serially). You can also change variable `defaultAs: "192.168.250.1"` to specify the target of ICMP Echo Request packets sent by the emulated UEs. Selecting the IP address of a real-world server (e.g., `8.8.8.8`) is a good test of end-to-end connectivity. Finally, you can change the amount of information gNBsim outputs by modifying `logLevel` in the `logger` block at the end of the file. For any changes you make, just rerun `make aether-gnbsim-run` to see the effects; you do not need to reinstall gNBsim.

Clean Up

We recommend continuing on to the next section before wrapping

up, but when you are ready to tear down your Quick Start deployment of Aether, simply execute the following commands:

```
$ make aether-gnbsim-uninstall  
$ make aether-5gc-uninstall  
$ make aether-k8s-uninstall
```

Note that while we stepped through the system one component at a time, OnRamp includes compound Make targets. For example, you can uninstall everything covered in this section by typing:

```
$ make aether-uninstall
```

Look at the `Makefile` to see the available set of Make targets.

[1]

pod

Closer Look

Before tearing down your Quick Start deployment, there are three additional steps you can take to watch Aether in action. The first is to bring up the Aether Management Plane (AMP), which includes Dashboards showing different aspects of Aether's runtime behavior. The second is to inspect the logs written by the various microservices. The third is to enable packet capture, and then run an analysis tool to trace the flow of packets into and out of SD-Core.

Install AMP

The Aether Management Platform (AMP) is implemented by two Kubernetes applications: *Runtime Operational Control (ROC)* and a *Monitoring Service*.[\[1\]](#) AMP can be deployed on the same cluster as SD-Core by executing the following Make target:

```
$ make aether-amp-install
```

Once complete, `kubectl` will show the `aether-roc` and `cattle-monitoring-system` namespaces running in support of these two services, respectively, plus new `atomix` pods in the `kube-system` namespace. Atomix is the scalable key-value store that keeps the ROC data model persistent.

Note that what the implementation calls ROC, [Chapter 6](#) refers to generically as *Service Orchestration*.

You can access the dashboards for the two subsystems, respectively, at

```
http://<server_ip>:31194  
http://<server_ip>:30950
```

The programmatic API underlying the Control Dashboard, which was introduced in [Section 6.4](#), can be accessed at

`http://10.76.28.113:31194/aether-roc-api/` in our example deployment (where Aether runs on host `10.76.28.113`).

There is much more to say about the ROC and the Aether API, which we return to in the section on Runtime Control. For now, we suggest you simply peruse the Control Dashboard by starting with the dropdown menu in the upper right corner. For example, selecting *Devices* will show the set of UEs registered with Aether, similar to the screenshot in [Figure 48](#). In an operational setting, these values would be entered into the ROC through either the GUI or the underlying API. For the Quick Start scenario we're limiting ourselves to in this section, these values are loaded from `deps/amp/5g-roc/templates/roc-5g-models.json`.

Figure 48. Screenshot of the ROC dashboard, showing known *Devices*. The dropdown menu on the right lists other available pages.

Turning to the Monitoring Dashboard, you will initially see Kubernetes-related performance stats. Select the *5G Dashboard* option to display information reported by SD-Core. Similar to [Figure 49](#), the page shows an active (green) UPF, and once you rerun the RAN simulator (gNBsim), some number of active base stations and connected devices. The bottom panel shows the UPF throughput, which due to gNBsim's focus on stressing the control plane, typically shows only minimal activity.

Figure 49. Screenshot of the monitoring subsystem's 5G dashboard.

When you are done experimenting with AMP, type the following to tear it down:

```
$ make aether-amp-uninstall
```

View Logs

You've already seen the log file generated by gNBsim for each emulation run, but you can also inspect the logs generated by individual microservices that implement Aether. Doing so is certainly helpful when debugging a failure, but it can also be an aid

in learning how each microservice works. For example, the following command outputs the log for the `bessd` container, one of five containers running as part of the `upf-0` pod:

```
$ kubectl logs -n omece -p upf-0 bessd
```

While `kubectl` works just fine for tasks like this, you may also want to install [k9s](#), a terminal-based UI that provides a convenient alternative for interacting with Kubernetes. Once installed, the following command brings up the UI for the OMEC namespace that implements SD-Core.

```
$ k9s -n omece
```

[Figure 50](#) shows an example `k9s` display, where you can scroll up and down, and then invoke one of the listed commands—such as `<l>` (display log) or `<s>` (open a shell)—for the selected pod.

Figure 50. Screenshot of `k9s`'s UI for the OMEC namespace, with the AMF pod currently selected.

Run Ksniff and Wireshark

In addition to the trace output generated by the simulator, a good way to understand the inner working of Aether is to use [Ksniff](#) (a Kubernetes plugin) to capture packets and display their headers as they flow into and out of the microservices that implement Aether. Output from `Ksniff` can then be fed into [Wireshark](#).

To install the `Ksniff` plugin on the server running Aether, you need to first install `krew`, the Kubernetes plugin manager. Instructions on doing that can be found [online](#). Once that's done, you can install `Ksniff` by typing:

```
$ kubectl krew install sniff
```

You can then run `Ksniff` in the context of a specific Kubernetes pod by specifying their namespace and instance names, and then redirecting the output to `Wireshark`. If you don't have a desktop environment on your Aether server, you can either view the output

using a simpler packet analyzer, such as [tshark](#), or by redirecting the PCAP output in a file and transfer it a desktop machine for viewing in Wireshark.

For example, the following captures and displays traffic into and out of the AMF, where you need to substitute the name of the AMP pod you learned from `kubectl` in place of `amf-5887bbf6c5-pc9g2`.

```
$ kubectl sniff -n omecc amf-5887bbf6c5-pc9g2 -o - | tsha
```

Of course, you'll also need to restart the RAN emulator to generate workload for this tool to capture.

Scale Cluster

Everything up to this point has been done as part of the Quick Start configuration, with all the components running in a single server (VM or physical machine). We now describe how to scale Aether to run on multiple servers, where we assume this cluster-based configuration throughout the rest of this guide. Before continuing, though, you need to remove the Quick Start configuration by typing:

```
$ make aether-uninstall
```

There are two aspects of our deployment that scale independently. One is Aether proper: a Kubernetes cluster running the set of microservices that implement SD-Core and AMP (and optionally, other edge apps). The second is gNBsim: the emulated RAN that generates traffic directed at the Aether cluster. Minimally, two servers are required—one for the Aether cluster and one for gNBsim—with each able to scale independently. For example, having four servers would support a 3-node Aether cluster and a 1-node workload generator. This example configuration corresponds to the following `hosts.ini` file:

```
[all]
node1 ansible_host=172.16.144.50 ansible_user=aether ansible_port=22
node2 ansible_host=172.16.144.71 ansible_user=aether ansible_port=22
node3 ansible_host=172.16.144.18 ansible_user=aether ansible_port=22
node4 ansible_host=172.16.144.93 ansible_user=aether ansible_port=22

[master_nodes]
node1

[worker_nodes]
```

```
node2
node3
node4
```

```
[gnbsim_nodes]
node4
```

The first block identifies all the nodes; the second block designates which node runs the Ansible client and the Kubernetes control plane (this is the node you ssh into and invoke Make targets and `kubectl` commands); the third block designates the worker nodes being managed by the Ansible client; and the last block indicate which nodes run the gNBsim workload generator (gNBsim scales across multiple Docker containers, but these containers are **not** managed by Kubernetes). Note that having `master_nodes` and `gnbsim_nodes` contain exactly one common server (as we did previously) is what triggers Ansible to instantiate the Quick Start configuration.

You need to modify `hosts.ini` to match your target deployment. Once you've done that (and assuming you deleted your earlier Quick Start configuration), you can re-execute the same set of targets you ran before:

```
$ make aether-k8s-install
$ make aether-5gc-install
$ make aether-amp-install
$ make aether-gnbsim-install
$ make aether-gnbsim-run
```

This will run the same gNBsim test case as before, but originating in a separate VM. We will return to options for scaling up the gNBsim workload in a later section, along with describing how to run physical gNBs in place of gNBsim. Note that if you are primarily interested in the latter, you can still run Aether on a single server, and then connect that node to one or more physical gNBs.

Finally, apart from being able to run SD-Core and gNBsim on separate nodes—thereby cleanly decoupling the Core from the RAN—one question we have not yet answered is why you might want to

scale the Aether cluster to multiple nodes. One answer is that you are concerned about availability, so want to introduce redundancy.

A second answer is that you want to run some other edge application, such as an IoT or AI/ML platform, on the Aether cluster. Such applications can be co-located with SD-Core, with the latter providing local breakout. For example, OpenVINO is a framework for deploying inference models to process local video streams streams, for example, detecting and counting people who enter the field of view for 5G-connected cameras. Just like SD-Core, OpenVINO is deployed as a set of Kubernetes pods.

Further Reading

[OpenVINO Toolkit.](#)

A third possible answer is that you want to scale SD-Core itself, in support of a scalable number of UEs. For example, providing predictable, low-latency support for hundreds or thousands of IoT devices requires horizontally scaling the AMF. OnRamp provides a way to experiment with exactly that possibility. If you edit the core section of `vars/main.yml` to use an alternative values file (in place of `sdcore-5g-values.yml`):


```
values_file: "deps/5gc/roles/core/templates/hpa-5g-values"
```

you can deploy SD-Core with *Horizontal Pod Autoscaling (HPA)* enabled. Note that HPA is an experimental feature of SD-Core; it has not yet been officially released and is not yet supported.

Further Reading

[Horizontal Pod Autoscaling.](#)

[image] [image]



Verify Network

This section goes into depth on how SD-Core (which runs *inside* the Kubernetes cluster) connects to either physical gNBs or an emulated RAN (both running *outside* the Kubernetes cluster). For the purpose of this section, we assume you already have a scalable cluster running (as outlined in the previous section), SD-Core has been installed on that cluster, and you have a terminal window open on the Master node in that cluster.

Figure 51 shows a high-level schematic of Aether's end-to-end User Plane connectivity, where we start by focusing on the basics: a single Aether node, a single physical gNB, and just the UPF

container running inside SD-Core. The identifiers shown in gray in the figure (10.76.28.187, 10.76.28.113, ens18) are taken from our running example of an actual deployment (meaning your details will be different). All the other names and addresses are part of a standard Aether configuration.

Figure 51. The UPF pod running inside the server hosting Aether, with `core` and `access` bridging the two. Identifiers 10.76.28.187, 10.76.28.113, ens18 are specific to a particular deployment site.

As shown in the figure, there are two bridges that connect the physical interface (ens18 in our example) with the UPF container. The `access` bridge connects the UPF downstream to the RAN (this corresponds to 3GPP's N3 interface) and is assigned IP subnet 192.168.252.0/24. The `core` bridge connects the UPF upstream to the Internet (this corresponds to 3GPP's N6 interface) and is assigned IP subnet 192.168.250.0/24. This means, for example, that the `access` interface *inside* the UPF (which is assigned address 192.168.252.3) is the destination IP address of GTP-encapsulated user plane packets from the gNB.

Following this basic schematic, it is possible to verify that the UPF is connected to the network by checking to see that the `core` and `access` are properly configured. This can be done using `ip`, and you should see results similar to the following:

```
$ ip addr show core
15: core@ens3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500
    link/ether 06:f7:7c:65:31:fc brd ff:ff:ff:ff:ff:ff
    inet 192.168.250.1/24 brd 192.168.250.255 scope globa
        valid_lft forever preferred_lft forever
    inet6 fe80::4f7:7cff:fe65:31fc/64 scope link
        valid_lft forever preferred_lft forever

$ ip addr show access
14: access@ens3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500
    link/ether 82:ef:d3:bb:d3:74 brd ff:ff:ff:ff:ff:ff
    inet 192.168.252.1/24 brd 192.168.252.255 scope globa
        valid_lft forever preferred_lft forever
```

```

    inet6 fe80::80ef:d3ff:febb:d374/64 scope link
    valid_lft forever preferred_lft forever

```

The above output from `ip` shows the two interfaces visible to the server, but running *outside* the container. `kubect1` can be used to see what's running *inside* the UPF, where `bessd` is the name of the container image that implements the UPF, and `access` and `core` are the last two interfaces shown below:

```

$ kubect1 -n omec exec -ti upf-0 bessd -- ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UP
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
    valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
    valid_lft forever preferred_lft forever
3: eth0@if30: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1450 qdisc noqueue state UP
    link/ether 8a:e2:64:10:4e:be brd ff:ff:ff:ff:ff:ff
    inet 192.168.84.19/32 scope global eth0
    valid_lft forever preferred_lft forever
    inet6 fe80::88e2:64ff:fe10:4ebe/64 scope link
    valid_lft forever preferred_lft forever
4: access@if2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
    link/ether 82:b4:ea:00:50:3e brd ff:ff:ff:ff:ff:ff
    inet 192.168.252.3/24 brd 192.168.252.255 scope global access
    valid_lft forever preferred_lft forever
    inet6 fe80::80b4:eaff:fe00:503e/64 scope link
    valid_lft forever preferred_lft forever
5: core@if2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP
    link/ether 4e:ac:69:31:a3:88 brd ff:ff:ff:ff:ff:ff
    inet 192.168.250.3/24 brd 192.168.250.255 scope global core
    valid_lft forever preferred_lft forever
    inet6 fe80::4cac:69ff:fe31:a388/64 scope link
    valid_lft forever preferred_lft forever

```

When packets flowing upstream from the gNB arrive on the server's physical interface, they need to be forwarded over the `access` interface. This is done by having the following kernel route installed, which should be the case if your Aether installation was

successful:

```
$ route -n | grep "Iface\|access"
```

Destination	Gateway	Genmask	Flags	Me
192.168.252.0	0.0.0.0	255.255.255.0	U	0

Within the UPF, the correct behavior is to forward packets between the `access` and `core` interfaces. Upstream packets arriving on the `access` interface have their GTP headers removed and the raw IP packets are forwarded to the `core` interface. The routes inside the UPF's `bessd` container will look something like this:

```
$ kubectl -n ome exec -ti upf-0 -c bessd -- ip route
default via 169.254.1.1 dev eth0
default via 192.168.250.1 dev core metric 110
10.76.28.0/24 via 192.168.252.1 dev access
10.76.28.113 via 169.254.1.1 dev eth0
169.254.1.1 dev eth0 scope link
192.168.250.0/24 dev core proto kernel scope link src 19
192.168.252.0/24 dev access proto kernel scope link src
```

The default route via `192.168.250.1` directs upstream packets to the Internet via the `core` interface, with a next hop of the `core` interface outside the UPF. These packets then undergo source NAT in the kernel and are sent to the IP destination in the packet. This means that the `172.250.0.0/16` addresses assigned to UEs are not visible beyond the Aether server. The return (downstream) packets undergo reverse NAT and now have a destination IP address of the UE. They are forwarded by the kernel to the `core` interface by these rules on the server:

```
$ route -n | grep "Iface\|core"
```

Destination	Gateway	Genmask	Flags	Me
172.250.0.0	192.168.250.3	255.255.0.0	UG	0
192.168.250.0	0.0.0.0	255.255.255.0	U	0

The first rule above matches packets to the UEs on the `172.250.0.0/16` subnet. The next hop for these packets is the `core` IP address inside the UPF. The second rule says that next hop address is reachable on the `core` interface outside the UPF. As a

result, the downstream packets arrive in the UPF where they are GTP-encapsulated with the IP address of the gNB.

Note that if you are not finding `access` and `core` interfaces outside the UPF, the following commands can be used to create these two interfaces manually (again using our running example for the physical ethernet interface):

```
$ ip link add core link ens18 type macvlan mode bridge 1
$ ip link add access link ens18 type macvlan mode bridge
```

Beyond this basic understanding, there are three other details of note. First, we have been focusing on the User Plane because Control Plane connectivity is much simpler: RAN elements (whether they are physical gNBs or gNBsim) reach the AMF using the server's actual IP address (10.76.28.113 in our running example). Kubernetes is configured to forward SCTP packets arriving on port 38412 to the AMF container.

Second, the basic end-to-end schematic shown in [Figure 51](#) assumes each gNB is assigned an address on the same L2 network as the Aether cluster (e.g., 10.76.28.0/24 in our example scenario). This works when the gNB is physical or when we want to run a single gNBsim traffic source, but once we scale up the gNBsim by co-locating multiple containers on a single server, we need to introduce another network so each container has a unique IP address (even though the containers are all hosted on the same node). This more complex configuration is depicted in [Figure 52](#), where 172.20.0.0/16 is the IP subnet for the virtual network (also implemented by a Macvlan bridge, and named `gnbaccess`).

Figure 52. A server running multiple instances of gNBsim, connected to Aether.

For completeness, [Figure 53](#) shows the Macvlan setup for the Quick Start configuration, where both the `gnbaccess` bridge and gNBsim container run in the same server as the Core (but with the container managed by Docker, independent of Kubernetes).

Figure 53. The Quick Start configuration with all components running in a single server.

Finally, all of the configurable parameters used throughout this section are defined in the `core` and `gnbsim` sections of the `vars/main.yml` file. Note that an empty value for `core.ran_subnet` implies the physical L2 network is used to connect RAN elements to the core, as is typically the case when connecting physical gNBs.

```
core:
  standalone: "true"
  data_iface: ens18
  values_file: "config/sdcore-5g-values.yaml"
  ran_subnet: "172.20.0.0/16"
  helm:
    chart_ref: aether/sd-core
    chart_version: 0.12.6
  upf:
    ip_prefix: "192.168.252.0/24"
  amf:
    ip: "10.76.28.113"

gnbsim:
  ...
  router:
    data_iface: ens18
    macvlan:
      iface: gnbaccess
      subnet_prefix: "172.20"
```

Emulated RAN

gNBsim emulates a 5G RAN, generating (mostly) Control Plane traffic that can be directed at SD-Core. This section describes how to configure gNBsim to customize and scale the workload it generates. We assume gNBsim runs in one or more servers, independent of the server(s) that host SD-Core. These servers are specified in the `hosts.ini` file, as described in the section on Scaling Aether. This blueprint assumes you start with a variant of `vars/main.yml` customized for running gNBsim. This is easy to do:

```
$ cd vars
$ cp main-gnbsim.yml main.yml
```

Configure gNBsim

Two sets of parameters control gNBsim. The first set, found in the `gnbsim` section of `vars/main.yml`, controls how gNBsim is deployed: (1) the number of servers it runs on; (2) the number of Docker containers running within each server; (3) what configuration to run in each of those containers; and (4) how those containers connect to SD-Core. For example, consider the following variable definitions:

```
gnbsim:
  docker:
    container:
      image: omecproject/5gc-gnbsim:main-PR_88-cc0
      prefix: gnbsim
      count: 2
    network:
```

```

        macvlan:
            name: gnbnet

router:
    data_iface: ens18
    macvlan:
        iface: gnbaccess
        subnet_prefix: "172.20"

servers:
    0:
        - "config/gnbsim-s1-p1.yaml"
        - "config/gnbsim-s1-p2.yaml"
    1:
        - "config/gnbsim-s2-p1.yaml"
        - "config/gnbsim-s2-p2.yaml"

```

The `container.count` variable in the `docker` block specifies how many containers run in each server (2 in this example). The `router` block then gives the network specification needed for these containers to connect to the SD-Core; all of these variables are described in the previous section on Networking. Finally, the `servers` block names the configuration files that parameterize each container. In this example, there are two servers with two containers running in each, with `config/gnbsim-s2-p1.yaml` parameterizing the first container on the second server.

These config files then specify the second set of gNBsim parameters. A detailed description of these parameters is outside the scope of this guide (see <https://github.com/omec-project/gnbsim> for details), but at a high-level, gNBsim defines a set of *profiles*, each of which exercises a common usage scenario that the Core has to deal with. Each of these sequences is represented by a `profileType` in the config file. gNBsim supports seven profiles, which we list here:

```

- profileType: register                # UE Registration
- profileType: pdusessetest            # UE Initiated Sess
- profileType: anrelease               # Access Network (A
- profileType: uetrigger servicereq    # UE Initiated Serv
- profileType: deregister              # UE Initiated De-r

```

```
- profileType: nwtriggeruedereg      # Network Initiated
- profileType: uereqpdusessrelease   # UE Initiated Sess
```

The second profile (pdusettest) is selected by default. It causes the specified number of UEs to register with the Core, initiate a user plane session, and then send a minimal data packet over that session. Note that the rest of the per-profile parameters are highly redundant. For example, they specify the IMSI- and PLMD-related information UEs need to connect to the Core.

Finally, it is necessary to edit the `core` section of `vars/main.yml` to indicate the address at which gNBsim can find the AMF. For our running example, this would look like the following:

```
core:
  amf: "10.76.28.113"
```

Install/Uninstall gNBsim

Once you have edited the parameters (and assuming you already have SD-Core running), you are ready to install gNBsim. This includes starting up all the containers and configuring the network so they can reach the Core. This is done from the main OnRamp server you've been using, where you type:

```
$ make gnbsim-docker-install
$ make aether-gnbsim-install
```

Note that the first step may not be necessary, depending on whether Docker is already installed on the server(s) you've designated to host gNBsim.

When you are finished, the following uninstalls everything:

```
$ make aether-gnbsim-uninstall
```

Run gNBsim

Once gNBsim is installed and the Docker containers instantiated, you can run the simulation by typing

```
$ make aether-gnbsim-run
```

This can be done multiple times without reinstalling. For each run, you can use Docker to view the results, which have been saved in each of the containers. To do so, ssh into one of the servers designated to to run gNBsim, and then type

```
$ docker exec -it gnbsim-1 cat summary.log
```

Note that container name `gnbsim-1` is constructed from the prefix variable defined in the `docker` section of `vars/main.yml`, with `-1` indicating the first container.

In addition to scaling up the workload you put on the Core, you can also experiment with the emulation settings defined in any or all of the config files in `deps/gnbsim/config/`. Focusing on `profile2` in particular (because it sends data packets after registering each UE), variable `defaultAs: "192.168.250.1"` specifies the target of ICMP Echo Request packets. Changing the value to the IP address of a real-world server (e.g., `8.8.8.8`) causes the emulated UE to actually ping that server. Success is a good indication that your Aether cluster is properly configured to support end-to-end connectivity.

3:08

5G

←

🔍 ⋮

Edit access point

Name

TEST SIM

APN

internet

Proxy

Not set

Port

Not set

Username

Not set

Password

Not set

Server

Not set

MMSC

Not set

MMS proxy

Not set

3:09

5G

← Edit access point

🔍 ⋮

MMS proxy

Not set

MMS port

Not set

MCC

001

MNC

01

Authentication type

None

APN type

default,supl,dun,hipri,fota

APN protocol

IPv4/IPv6

APN roaming protocol

IPv4/IPv6


APN enable/disable

APN enabled

☒

Bearer

Unspecified



Physical RAN

We are now ready to replace the emulated RAN with physical gNBs and real UEs. You will need to edit `hosts.ini` to reflect the Aether cluster you want to support, where just a single server is sufficient and there is no reason to include nodes in the `[gnbsim_nodes]` set.

In addition to the server(s) in your cluster, bringing up a physical RAN requires the following hardware:

- One or more 5G small cell radios (e.g., MOSO CANOPY 5G INDOOR SMALL CELL).
- One or more 5G-capable UEs (e.g., unlocked Motorola Moto G 5G).
- A SIM reader/writer and associated software (e.g., OYEITIMES MCR3516).

- A set of programmable SIM cards (blank cards likely included with reader/writer).

There are multiple options for each component, but finding a combination that works together can be challenging. This section makes several recommendations based on working end-to-end systems. For simplicity, we pared the above list back to a single example for each item, but these should not be interpreted as the only possibility.

Note also that our example relies on the availability of spectrum in the CBRS band, which is available in the United States. Spectrum options are likely to be different in other countries.

Troubleshooting Hint

We are tracking community experience with different hardware in the `#aether-onramp` channel of the [ONF Workspace](#), with summaries of different combinations people have tried reported in the OnRamp [Troubleshooting Wiki Page](#).

This blueprint assumes you start with a variant of `vars/main.yml` customized for running physical 5G radios. This is easy to do:

```
$ cd vars
$ cp main-gNB.yml main.yml
```

This section focuses on bringing up a single gNB, and assumes that gNB is on the same L2 network as the Aether cluster. In our running example, this implies both are on subnet `10.76.28.0/24`.

Troubleshooting Hint

Be aware that enterprise and campus networks have been known to filter packets in ways that impact gNB-to-Core connectivity. Fortunately, physical gNBs connect to SD-Core (both the AMF in the control plane and the UPF in the user plane) in exactly the same way as external instances of gNBsim. Going through the process of bringing up gNBsim in a second server, as described in the previous section, is a good way to validate that your Core is “gNB-ready”.

Modify Configuration

Modify the `core` section of `vars/main.yml` to match the following, substituting your local details for `ens18` and `10.76.28.113`. Of particular note, setting `ran_subnet` to the empty string indicates that the gNB is connected to the same physical L2 network as your Aether cluster, and the new `values_file` is tailored for a physical RAN rather than the emulated RAN we've been using.

```
core:
  standalone: "true"
  data_iface: ens18
  values_file: "deps/5gc/roles/core/templates/radio-5g"
  ran_subnet: ""
  helm:
    chart_ref: aether/sd-core
    chart_version: 0.12.6
  upf:
    ip_prefix: "192.168.252.0/24"
  amf:
    ip: "10.76.28.113"
```

Prepare UEs

When selecting UEs to connect to Aether, be aware that not all phones support the CBRS frequency bands that Aether uses. For band n78, Aether is known to work with recent iPhones (11 and greater), Google Pixel phones (4 and greater), OnePlus phones, and Moto G 5G phones. For band n48, Aether is known to work with Moto G 5G and OnePlus phones; Pixel 7 phones are purported to work as well. Another option is to use a 5G dongle connected to a Raspberry Pi as a demonstration UE. This makes it easier to run diagnostic tests from the UE. For example, we have used [APAL's 5G dongle](#) with Aether.

5G-connected devices must have a SIM card, which you are responsible for programming and inserting. You will need a SIM

card writer, such as the *OYEITIMES MCR3516* (available on Amazon), which comes with five blank cards. You will need to set the PLMN identifier (constructed from a valid MCC/MNC pair), the IMSI, and two secret keys for each SIM card. As working examples, we have used two different PLMN ids: 315010 constructed from MCC=315 (US) and MNC=010 (CBRS), and 00101 constructed from MCC=001 (TEST) and MNC=01 (TEST). You should use whatever values are appropriate for your local environment, where we use the following as a running example throughout this section:

- IMSI: each one is unique, matching pattern
315010***** (up to 15 digits)
- OPc: 69d5c2eb2e2e624750541d3bbc692ba5
- Key: 000102030405060708090a0b0c0d0e0f

Note that the actual config files distributed with OnRamp have IMSIs constructed using PLMN id 00101. Both sets of examples are taken from working deployments (315010 for a 4G/eNB and 00101 for a 5G/gNB), so in principle either should work as a model you can emulate in your deployment. As a practical matter, however, it is certainly easiest (and safest) to start with the existing code.

After inserting the SIM card into the device and powering it up, log into the phone, select *Network Settings* > *SIMs* and create a new *Access Point Name (APN)*, configured as shown in [Figure 54](#). The entry name (TEST SIM in the example) is arbitrary and the MCC/MNC pair is set automatically based on the newly inserted SIM card. The important value is the APN, which is set to *internet*. This value corresponds to variable *dnn (Data Network Name)* defined in `deps/5gc/roles/core/templates/radio-5g-values.yaml`. Loosely speaking, the role the APN plays in the mobile network is similar to the role an SSID plays in a WiFi network.

Figure 54. Configure an Access Point Name (APN) for the new SIM card on the UE.

Finally, modify the `subscribers` block of the `omec-sub-provision` section in file `deps/5gc/roles/core/templates/radio-5g-values.yaml` to record the IMSI, OPc, and Key values

configured onto your SIM cards. The block also defines a sequence number that is intended to thwart replay attacks. For example, the following code block adds IMSIs between 315010999912301 and 315010999912310:

```
subscribers:
- ueId-start: "315010999912301"
  ueId-end: "315010999912310"
  plmnId: "315010"
  opc: "69d5c2eb2e2e624750541d3bbc692ba5"
  key: "000102030405060708090a0b0c0d0e0f"
  sequenceNumber: 135
```

Further down in the same `omec-sub-provision` section you will find two other blocks that also need to be edited. The first, `device-groups`, assigns IMSIs to *Device Groups*. You will need to reenter the individual IMSIs from the `subscribers` block that will be part of the device-group:

```
device-groups:
- name: "5g-user-group1"
  imsis:
    - "315010999912301"
    - "315010999912302"
    - "315010999912303"
```

The second block, `network-slices`, sets various parameters associated with the *Slices* that connect device groups to applications. Here, you will need to reenter the PLMN information, with the other slice parameters remaining unchanged (for now):

```
plmn:
  mcc: "315"
  mnc: "010"
```

Aether supports multiple *Device Groups* and *Slices*, but the data entered here is purposely minimal; it's just enough to bring up and debug the installation. Over the lifetime of a running system, information about *Device Groups* and *Slices* (and the other abstractions they build upon) should be entered via the ROC, as

described the section on Runtime Control. When you get to that point, Ansible variable `standalone` in `vars/main.yml` (which corresponds to the override value assigned to `provision-network-slice` in `radio-5g-values.yml`) should be set to `false`. Doing so causes the `device-groups` and `network-slices` blocks of `radio-5g-values.yml` to be ignored. The `subscribers` block is always required to configure SD-Core.

Bring Up Aether

You are now ready to bring Aether on-line. We assume a fresh install by typing the following:

```
$ make aether-k8s-install
$ make aether-5gc-install
```

You can verify the installation by running `kubectl` just as you did in earlier stages. Note that we postpone bringing up the AMP until later so as to have fewer moving parts to debug.

gNodeB Setup

Once the SD-Core is up and running, we are ready to bring up the physical gNB. The details of how to do this depend on the specific device you are using, but we identify the main issues you need to address using SERCOMM's 5G femto cell (as distributed by MosoLabs) as an example. That particular device uses either the n48 or n78 band and is on the ONF MarketPlace, where you can also find a User's Guide that gives detailed instructions about configuring the gNB.

Further Reading

[MOSO CANOPY 5G INDOOR SMALL CELL.](#)

Troubleshooting Hint

The product data sheet shows support for frequency bands n78/

n48/n77, but individual devices do not necessarily support all three. For example, we have experience with an n78 device and an n48 device, with the latter (n48) becoming the preferred band (due in part to less risk of interfering with Radio Altimeters). For n48, PLMN id 00101 is currently recommended.

For the purposes of the following description, we assume the gNB is assigned IP address 10.76.28.187, which per our running example, is on the same L2 network as our Aether server (10.76.28.113). [Figure 55](#) shows a screenshot of the SERCOMM gNB management dashboard, which we reference in the instructions that follow:

Figure 55. Management dashboard on the Sercomm gNB, showing the dropdown Settings menu overlayed on the NR Cell Configuration page (which shows default radio settings).

1. **Connect to Management Interface.** Start by connecting a laptop directly to the LAN port on the small cell, pointing your laptop's web browser at the device's management page (<https://10.10.10.189>). You will need to assign your laptop an IP address on the same subnet (e.g., 10.10.10.100). Once connected, log in with the credentials provided by the vendor.
2. **Configure WAN.** Visit the `Settings > WAN` page to configure how the small cell connects to the Internet via its WAN port, either dynamically using DHCP or statically by setting the device's IP address (10.76.28.187) and default gateway (10.76.28.1).
3. **Access Remote Management.** Once on the Internet, it should be possible to reach the management dashboard without being directly connected to the LAN port (<https://10.76.28.187>).
4. **Connect GPS.** Connect the small cell's GPS antenna to the GPS port, and place the antenna so it has line-of-sight to the sky (i.e., place it in a window). The `Status` page of the management dashboard should report its latitude, longitude, and fix time.
5. **Spectrum Access System.** One reason the radio needs GPS is so it can report its location to a Spectrum Access System

(SAS), a requirement in the US to coordinate access to the CBRS Spectrum in the 3.5 GHz band. For example, the production deployment of Aether uses the [Google SAS portal](#), which the small cell can be configured to query periodically. To do so, visit the `Settings > SAS` page. Acquiring the credentials needed to access the SAS requires you go through a certification process, but as a practical matter, it may be possible to test an isolated/low-power femto cell indoors before completing that process. Consult with your local network administrator.

6. **Configure Radio Parameters.** Visit the `Settings > NR Cell Configuration` page (shown in the figure) to set parameters that control the radio. It should be sufficient to use the default settings when getting started.
7. **Configure the PLMN.** Visit the `Settings > 5GC` page to set the PLMN identifier on the small cell (00101) to match the MCC/MNC values (001 / 01) specified in the Core.
8. **Connect to Aether Control Plane.** Also on the `Settings > 5GC` page, define the AMF Address to be the IP address of your Aether server (e.g., 10.76.28.113). Aether's SD-Core is configured to expose the corresponding AMF via a well-known port, so the server's IP address is sufficient to establish connectivity. The `Status` page of the management dashboard should confirm that control interface is established.
9. **Connect to Aether User Plane.** As described in an earlier section, the Aether User Plane (UPF) is running at IP address 192.168.252.3. Connecting to that address requires installing a route to subnet 192.168.252.0/24. How you install this route is device and site-dependent. If the small cell provides a means to install static routes, then a route to destination 192.168.252.0/24 via gateway 10.76.28.113 (the server hosting Aether) will work. If the small cell does not allow static routes (as is the case for the SERCOMM gNB), then 10.76.28.113 can be installed as the default gateway, but doing so requires that your server also be configured to forward IP packets on to the Internet.

Troubleshooting Hint

For the SERCOMM gNB, if you elect to enable GPS, then `Setting > Sync_Settings > Sync_Mode` should be set to `TIME`. With GPS and PTP disabled, `Setting > Sync_Settings > Sync_Mode` should be set to `FREE_RUNNING`.

Troubleshooting Hint

For the SERCOMM gNB, we recommend the following when the gNB's addresses is acquired via DHCP, assuming that address is unlikely to change. When configuring the WAN (via the LAN), start with DHCP enabled. Note the IP address the gNB has been assigned, and then after disconnecting from the LAN, connect to the GUI via this address. You will be on the same L2 subnet as the Aether server, which you should be able to ping using the gNB's diagnostic tool. The default gateway DHCP returns does not know how to route data packets to the UPF. To fix this, modify the WAN settings to use a static IP, with the DHCP-provided IP used as the gNB's static address. Then set the default gateway to the IP address of your Aether server.

Run Diagnostics

Successfully connecting a UE to the Internet is not a straightforward exercise. It involves configuring the UE, gNB, and SD-Core software in a consistent way; establishing SCTP-based control plane (N2) and GTP-based user plane (N3) connections between the base station and Mobile Core; and traversing multiple IP subnets along the end-to-end path.

The UE and gNB provide limited diagnostic tools. For example, it's possible to run `ping` and `traceroute` from both. You can also run the `ksniff` tool described in the Networking section, but the most helpful packet traces you can capture are shown in the following commands. You can run these on the Aether server, where we use our example `ens18` interface for illustrative purposes:

```
$ sudo tcpdump -i any sctp -w sctp-test.pcap
$ sudo tcpdump -i ens18 port 2152 -w gtp-outside.pcap
```

```
$ sudo tcpdump -i access port 2152 -w gtp-inside.pcap
$ sudo tcpdump -i core net 172.250.0.0/16 -w n6-inside.p
$ sudo tcpdump -i ens18 net 172.250.0.0/16 -w n6-outside
```

The first trace, saved in file `sctp.pcap`, captures SCTP packets sent to establish the control path between the base station and the Mobile Core (i.e., N2 messages). Toggling “Mobile Data” on the UE, for example by turning Airplane Mode off and on, will generate the relevant control plane traffic.

The second and third traces, saved in files `gtp-outside.pcap` and `gtp-inside.pcap`, respectively, capture GTP packets (tunneled through port 2152) on the RAN side of the UPF. Setting the interface to `ens18` corresponds to “outside” the UPF and setting the interface to `access` corresponds to “inside” the UPF. Running `ping` from the UE will generate the relevant user plane (N3) traffic.

Similarly, the fourth and fifth traces, saved in files `n6-inside.pcap` and `n6-outside.pcap`, respectively, capture IP packets on the Internet side of the UPF (which is known as the N6 interface in 3GPP). In these two tests, `net 172.250.0.0/16` corresponds to the IP addresses assigned to UEs by the SMF. Running `ping` from the UE will generate the relevant user plane traffic.

If the `gtp-outside.pcap` has packets and the `gtp-inside.pcap` is empty (no packets captured), you may run the following commands to make sure packets are forwarded from the `ens18` interface to the `access` interface and vice versa:

```
$ sudo iptables -A FORWARD -i ens18 -o access -j ACCEPT
$ sudo iptables -A FORWARD -i access -o ens18 -j ACCEPT
```

Support for eNBs

Aether OnRamp is geared towards 5G, but it does support physical eNBs, including 4G-based versions of both SD-Core and AMP. It does not support an emulated 4G RAN. The 4G blueprint uses all

the same Ansible machinery outlined in earlier sections, but starts with a variant of `vars/main.yml` customized for running physical 4G radios:

```
$ cd vars
$ cp main-eNB.yml main.yml
```

Assuming that starting point, the following outlines the key differences from the 5G case:

1. There is a 4G-specific repo, which you can find in `deps/4gc`.
2. The `core` section of `vars/main.yml` specifies a 4G-specific values file:

```
values_file: "deps/4gc/roles/core/templates/
radio-4g-values.yaml"
```

3. The `amp` section of `vars/main.yml` specifies that 4G-specific models and dashboards get loaded into the ROC and Monitoring services, respectively:

```
roc_models: "deps/amp/roles/roc-load/
templates/roc-4g-models.json"
```

```
monitor_dashboard: "deps/amp/roles/monitor-
load/templates/4g-monitor"
```

4. You need to edit two files with details for the 4G SIM cards you use. One is the 4G-specific values file used to configure SD-Core:

```
deps/4gc/roles/core/templates/radio-4g-
values.yaml
```

The other is the 4G-specific Models file used to bootstrap ROC:

```
deps/amp/roles/roc-load/templates/radio-4g-
models.json
```

5. There are 4G-specific Make targets for SD-Core (e.g., `make aether-4gc-install` and `make aether-4gc-uninstall`), but the Make targets for AMP (e.g., `make aether-amp-install` and `make aether-amp-uninstall`) work unchanged in both 4G and 5G.

The Quick Start and Emulated RAN (gNBsim) deployments are for 5G only, but revisiting the other sections—substituting the above for their 5G counterparts—serves as a guide for deploying a 4G version of Aether. Note that the network is configured in exactly the same way for both 4G and 5G. This is because SD-Core’s implementation of the UPF is used in both cases.

Runtime Control

Aether defines an API (and associated GUI) for managing connectivity at runtime. This stage brings up that API/GUI, as implemented by the *Runtime Operational Control (ROC)* subsystem, building on the physical gNB we connected to Aether in the previous section.

This stage focuses on the abstractions that the ROC layers on top of the SD-Core. These abstractions are described in [Section 6.4](#) and include *Device Groups* and *Slices*. (The full set of model definitions can be found in [GitHub](#).) Initial settings of these ROC-managed parameters are recorded in `deps/amp/roles/roc-load/templates/radio-5g-models.json`. We use these values to load the ROC database, saving us from a laborious GUI session.

Somewhat confusingly, the *Device-Group* and *Slice* information is duplicated between `deps/5gc/roles/core/templates/radio-5g-values.yaml` and this `radio-5g-models.json` file. This makes it possible to bring up the SD-Core without the ROC, which simplifies the process of debugging an initial installation, but having two sources for this information leads to problems keeping them in sync, and should be avoided.

To this end, Aether treats the ROC as the “single source of truth” for *Slices*, *Device Groups*, and all the other abstract objects it defines, so we recommend using the GUI or API to make changes over time, and avoiding the override values in `radio-5g-values.yaml` once you’ve established basic connectivity. And if you want to save this bootstrap state in a text file for a possible restart, we recommend doing so in `radio-5g-models.json` (although this is not a substitute for the operational practice of backing up the ROC database).

To make ROC the authoritative source of runtime state, first edit the standalone variable in the core section of vars/main.yml, setting it to false. This variable indicates whether we want SD-Core to run in *Stand Alone* mode, which has been the default setting up to this point. Disabling standalone causes the SD-Core to ignore the device-groups and network-slices blocks of the omec-sub-provision section in radio-5gc-values.yaml, and to instead retrieve this information from the ROC.

The next step is to edit radio-5g-models.json to record the same IMSI information you added to radio-5gc-values.yaml in the previous section. This includes modifying, adding and removing sim-card entries as necessary. Note that only the IMSIs need to match the earlier data; the sim-id and display-name values are arbitrary and need only be consistent *within* radio-5g-models.json.

```
"imsi-definition": {
    "mcc": "315",
    "mnc": "010",
    "enterprise": 1,
    "format": "CCCNNEESSSSSS"
},
...

"sim-card": [
    {
        "sim-id": "sim-1",
        "display-name": "SIM 1",
        "imsi": "315010999912301"
    },
    ...
]
```

Once you are done with these edits, uninstall the SD-Core you had running in the previous stage, and then bring up the ROC followed by a new instantiation of the SD-Core:

```
$ make aether-5gc-uninstall
$ make aether-amp-install
$ make aether-5gc-install
```

The order is important, since the Core depends on configuration parameters provided by the ROC. Also note that you may need to reboot the gNB, although it typically does so automatically when it detects that the Core has restarted.

To see these initial configuration values using the GUI, open the dashboard available at `http://<server-ip>:31194`. If you select `Configuration > Site` from the drop-down menu at top right, and click the `Edit` icon associated with the `Aether Site` you can see (and potentially change) the following values:

- MCC: 315
- MNC: 010

Although we have no need to do so now, you can make changes to these values, and then click `Update` to save them to the “commit basket”. Similarly, if you select `Sim Cards` from the drop-down menu at top right, the `Edit` icon associated with each SIM card allows you to see (and potentially change) the IMSI values associated with each device. You can also disable individual IMSIs. Again, click `Update` if you make any changes.

The set of registered IMISs can be aggregated into *Device-Groups* by selecting `Device-Groups` from the drop-down menu at the top right, and adding a new device group.

Finally, if you do make a set of updates, select the `Basket` icon at top right when you are done, and click the `Commit` button. This causes the set of changes to be committed as a single transaction.

A more complete User’s Guide for the ROC is available online, although be aware that our OnRamp-based deployment has not yet enabled the secure login feature.

Further Reading

[Aether Operations](#).